

PAPER • OPEN ACCESS

How to detect lensing rotation

To cite this article: Mathew Robertson and Antony Lewis JCAP08(2023)048

View the [article online](#) for updates and enhancements.

You may also like

- [CMB-S4: Forecasting Constraints on Primordial Gravitational Waves](#)
Kevork Abazajian, Graeme E. Addison, Peter Adshead et al.
- [Current and Future Constraints on Primordial Magnetic Fields](#)
Dylan R. Sutton, Chang Feng and Christian L. Reichardt
- [A Cross-internal Linear Combination Approach to Probe the Secondary CMB Anisotropies: Kinematic Sunyaev–Zel'dovich Effect and CMB Lensing](#)
Srinivasan Raghunathan and Yuuki Omori

How to detect lensing rotation

Mathew Robertson and Antony Lewis

Department of Physics and Astronomy, University of Sussex,
Brighton BN1 9QH, U.K.

E-mail: mr671@sussex.ac.uk, antony@cosmologist.info

Received March 28, 2023

Revised July 21, 2023

Accepted July 31, 2023

Published August 18, 2023

Abstract. Gravitational lensing rotation of images is predicted to be negligible at linear order in density perturbations, but can be produced by the post-Born lens-lens coupling at second order. This rotation is somewhat enhanced for Cosmic Microwave Background (CMB) lensing due to the large source path length, but remains small and very challenging to detect directly by CMB lensing reconstruction alone. We show the rotation may be detectable at high significance as a cross-correlation signal between the curl reconstructed with Simons Observatory (SO) or CMB-S4 data, and a template constructed from quadratic combinations of large-scale structure (LSS) tracers. Equivalently, the lensing rotation-tracer-tracer bispectrum can also be detected, where LSS tracers considered include the CMB lensing convergence, galaxy density, and the Cosmic Infrared Background (CIB), or optimal combinations thereof. We forecast that an optimal combination of these tracers can probe post-Born rotation at the level of 5.7σ – 6.1σ with SO and 13.6σ – 14.7σ for CMB-S4, depending on whether standard quadratic estimators or maximum a posteriori iterative methods are deployed. We also show possible improvement up to 21.3σ using a CMB-S4 deep patch observation with polarization-only iterative lensing reconstruction. However, these cross-correlation signals have non-zero bias because the rotation template is quadratic in the tracers, and exists even if the lensing is rotation free. We estimate this bias analytically, and test it using simple null-hypothesis simulations to confirm that the bias remains subdominant to the rotation signal of interest. Detection and then measurement of the lensing rotation cross-spectrum is therefore a realistic target for future observations.

Keywords: CMBR polarisation, galaxy clusters, gravitational lensing, weak gravitational lensing

ArXiv ePrint: [2303.13313](https://arxiv.org/abs/2303.13313)

Contents

1	Introduction	1
2	Background	2
2.1	Lensing convergence	3
2.1.1	Convergence power spectrum	3
2.2	Post-Born rotation	4
2.2.1	Rotation power spectrum	5
3	Is rotation detectable?	5
3.1	CMB lensing reconstruction	5
3.1.1	CMB experimental setups	6
3.2	CMB rotation forecasts	7
3.3	Galaxy lensing rotation	10
4	Optimal rotation template	10
4.1	Template construction	11
4.2	Tracers	12
4.2.1	Tracer choices	12
4.3	Fast real-space template construction	14
5	Forecasts	15
5.1	Bispectra forecasts	15
5.2	Cross-spectra forecasts	16
6	Cross-spectrum bias	19
6.1	Template-lensing bias	20
6.2	Additional biases with $\hat{\kappa}$	22
6.3	Testing using simulations	25
7	Conclusions	26
A	Iterative estimators	27
B	Density tracers	28
B.1	CMB lensing convergence	28
B.2	Galaxy clustering	28
B.3	Cosmic Infrared Background	29
C	Isolating cross-spectrum biases	31
C.1	$N_{A1}^{(2)} + N_{C1}^{(2)}$	31
C.2	$N_{\hat{\kappa}\hat{\kappa}}^{(0)}$	32
C.3	$N_{\hat{\kappa}}^{(1)}$	33

1 Introduction

Gravitational lensing of Cosmic Microwave Background (CMB) photons by the matter along our line of sight distorts the observed CMB fields, giving a distinct observational signature of lensing (see ref. [1] for a review). The physics generating the primordial anisotropies in the CMB fields are well understood, and they are very close to Gaussian, so the non-Gaussian statistics of the observed lensed fields contain enough information to reconstruct the lensing field up to statistical fluctuations. The power spectrum of the lensing field depends on the large-scale structure (LSS) density fluctuations and late-time geometry of the Universe, and hence provides an observational probe of the Universe after recombination. The CMB lensing convergence has now been detected to high significance by multiple experiments [2–8], corresponding to purely gradient lensing deflection angles as predicted for first-order lensing by density perturbations.

Higher-order corrections to the lensing field can however produce lensing rotation from scalar perturbations, as predicted by a second-order post-Born lensing expansion [9–13]. This corresponds to having a curl-like component to the lensing deflection field [14], and is produced by a lens-lens coupling between non-aligned shear distortions induced by two gravitational potentials at different redshifts. Since the rotation is second order, it is predicted to be much smaller than the convergence signal. For current experiments it can be taken to be negligible, so that the measured rotation signal can serve as a null test. The predicted lensing post-Born rotation can in principle be measured by lensing quadratic estimators with more sensitive future data; however, as we shall describe, even using iterative estimators with fourth-generation (CMB-S4) experiments, it is challenging to detect directly at any significance due to the large lensing reconstruction noise.

Ref. [12] showed that the rotation may instead be detectable with CMB-S4 via the convergence-convergence-rotation bispectrum, though more futuristic data would still be required to get much above a 3σ detection. Conceptually, this bispectrum estimator is using two convergence modes to probe the two lenses that combine to produce the rotation via lens-lens coupling, giving an approximate quadratic estimate of the lensing rotation. The cross-correlation of this quadratic rotation estimate with the measured rotation then gives a probe of the lensing rotation power. This immediately suggests that we could do much better: by using other tracers of the large-scale structure with lower noise and better redshift resolution, we can trace the multiple redshifts of the lens-lens couplings to produce a more accurate rotation template. In the limit in which we had a perfect tracer of the LSS mass distribution, it would be possible to predict exactly the lensing rotation that would be produced by ray tracing through it. This perfect rotation template could then be correlated with the lensing reconstruction rotation field, giving a much more sensitive probe of the rotation than the reconstruction auto-spectrum. In this paper we investigate how close we can get to this ideal situation, by combining more realistic large-scale probes (such as galaxy counts and Cosmic Infrared Background (CIB)) as well as the CMB lensing convergence.

We begin with an overview of weak lensing observables and modelling of their power spectra, including post-Born rotation, in section 2. We then examine the difficulty of direct lensing rotation reconstruction in section 3, but prove that detection is feasible for near-future CMB missions with some appropriately constructed rotation template. In section 4 we construct an explicit estimator for this lensing rotation template, and relate its cross-spectrum with the measured rotation to the corresponding bispectrum estimators. We then make forecasts in section 5 for the upcoming Simons Observatory (SO) [15] experiment and

also for the next generation “stage 4” (CMB-S4) [16] experiment, combined with various LSS tracers, showing that the rotation signal should be detectable at high significance. As a first-look analysis, we make fairly idealized assumptions, but show that even with idealized data there will be some non-zero biases affecting the rotation power spectrum estimate. Section 6 describes the origin of these biases, and we construct an explicit analytic model which we then test against simulations. Our new estimators and models provide a way to clearly detect the lensing rotation signal with future data, and can also be used to quantify the expected signal in mixed rotation bispectra that may be useful null tests for other cosmological analyses. We finish by summarizing such conclusions in section 7.

Throughout this paper we assume a standard flat Λ CDM cosmology consistent with Planck [17]. This only enters at the level of the matter power spectrum computed by CAMB [18], with non-linear corrections modelled by HMCODE [19].

2 Background

Weak lensing is the statistical distortion of an observation due to the gravitational influence of the intervening LSS. Photons are deflected from their background geodesic path by an angle proportional to the gradient of the gravitational potential $\nabla\Psi(\chi)$. The total deflection, α , is the accumulation of all gravitational effects inflicted upon a particular ray [1]

$$\alpha(\hat{n}) = -2 \int_0^{\chi_s} d\chi W(\chi, \chi_s) \nabla\Psi(\hat{n}, \chi), \quad (2.1)$$

where the window function, $W(\chi, \chi_s)$, is the lensing kernel for an image at source distance χ_s , and χ is the comoving radial distance.

The observed position, $\hat{n}(\theta, \phi)$, of an observable, X , on the celestial sphere has been transported from its true position by a deflection angle α , so that the measured *lensed* observable is actually a remapping of the unlensed field,¹

$$\tilde{X}(\hat{n}) = X(\hat{n} + \alpha(\hat{n})). \quad (2.2)$$

In eq. (2.1), at lowest order α is determined by integrating the contributions of $\nabla\Psi$ along the background geodesic. However, photons do not travel along background geodesics due to the lensing deflection. To account for these deviations, perturbative corrections to the angular position at which the gravitational potential is evaluated along the radial path must be applied within the integral. At some arbitrary point along the radial path at position \hat{n}_0 , we can expand about the background geodesic to second order [12]

$$\begin{aligned} \Psi(\hat{n}) &= \Psi(\hat{n}_0 + \alpha(\hat{n}_0)) \\ &= \Psi(\hat{n}_0) + [\alpha^a \nabla_a \Psi](\hat{n}_0) + \mathcal{O}(\Psi^3), \end{aligned} \quad (2.3)$$

where implicit dependence between the gravitational potential and deflection angle results in the iterative solution to eq. (2.1), which to second order is

$$\alpha_a(\hat{n}_0) = -2 \int_0^{\chi_s} d\chi W(\chi, \chi_s) \left[\nabla_a \Psi(\chi) - 2 \int_0^\chi d\chi' W(\chi', \chi) \nabla_a \nabla_b \Psi(\chi) \nabla^b \Psi(\chi') \right] (\hat{n}_0) + \mathcal{O}(\Psi^3). \quad (2.4)$$

Higher order corrections to α are relatively tame to compute, but here we are only interested in the leading-order contribution to the lensing rotation, which sits nicely at second order.

¹It is not exactly a remapping, but remapping is a good approximation [20].

2.1 Lensing convergence

It can be seen from eq. (2.4) (or eq. (2.1)) that to first order the total deflection is determined by the gravitational potential along the background geodesic. This is known as the Born approximation, which describes one effective total deflection from the original path of a photon. At this order the deflection field is curl free,² i.e. $\nabla \times \boldsymbol{\alpha} = 0$, and thus is fully described as the angular gradient of some function

$$\boldsymbol{\alpha}(\hat{\mathbf{n}}) = \nabla \phi(\hat{\mathbf{n}}), \quad (2.5)$$

where that function ϕ is known as the *lensing potential*. At this order ϕ is simply the projection of the gravitational potentials sourced from density perturbations along the unperturbed geodesic.

The *lensing convergence* is the observable describing the magnification effect and is directly related to the lensing potential by

$$\kappa(\hat{\mathbf{n}}) = -\frac{1}{2} \nabla \cdot \boldsymbol{\alpha}(\hat{\mathbf{n}}) = -\frac{1}{2} \nabla^2 \phi(\hat{\mathbf{n}}). \quad (2.6)$$

Working within the flat sky approximation of the two dimensional spherical sky, the convergence in harmonic space is written

$$\kappa(\mathbf{L}) = L^2 \int_0^{\chi_s} d\chi W(\chi, \chi_s) \Psi(\mathbf{L}, \chi). \quad (2.7)$$

For CMB lensing, the source is emitted at the surface of last scattering, $\chi_s = \chi_*$, which can be taken as a single source plane to high accuracy.

2.1.1 Convergence power spectrum

Using the usual definition for the power spectrum of two fields in harmonic space $\langle X(\mathbf{L})Y(\mathbf{L}') \rangle = (2\pi)^2 \delta_{\mathbf{D}}(\mathbf{L} + \mathbf{L}') C_L^{XY}$, and the lowest-order Limber approximation, $k \approx L/\chi$ [21], the convergence power spectrum to leading order is

$$C_L^{\kappa\kappa}(\chi_1, \chi_2) = \int_0^{\chi_2} d\chi \frac{W_\kappa(\chi, \chi_1) W_\kappa(\chi, \chi_2)}{\chi^2} P_{\delta\delta} \left(k \approx \frac{L}{\chi}, z(\chi) \right), \quad (2.8)$$

for two sources at radial distances χ_1 and χ_2 where $\chi_2 \geq \chi_1$. This distinction between source planes becomes necessary later in the paper when defining lensing bispectra.

Note that we have switched from integrating over the Weyl potential, Ψ , to the comoving-gauge matter density perturbation, δ , for later notational convenience. The conversion is achieved through the cosmological Poisson equation which has a simple form during matter and dark energy domination (assuming General Relativity and neglecting small effects from massive neutrinos):

$$k^2 \Psi(\mathbf{k}, \chi) \approx -\gamma(\chi) \delta(\mathbf{k}, \chi) = -\frac{3\Omega_m H_0^2}{2a(\chi)} \delta(\mathbf{k}, \chi), \quad (2.9)$$

for scale factor $a(\chi) = (1 + z(\chi))^{-1}$, present-day density parameter of matter, Ω_m , and Hubble constant, H_0 .

²As we are working on the two-dimensional sky, the cross product here is defined as $(\nabla \times A)_a = \epsilon_{ab} \nabla_b A$ where ϵ is the antisymmetric tensor or the 2-dimensional Levi-Cevita symbol.

The window function for the CMB lensing convergence is³

$$W_\kappa(\chi, \chi_s) = \gamma(\chi) \chi^2 \frac{\chi_s - \chi}{\chi \chi_s} \Theta(\chi_s - \chi), \quad (2.10)$$

where the Heaviside function Θ ensures $\chi_s \geq \chi$. For weak lensing with galaxies, the galaxy sources lie within many redshift shells, unlike the CMB which is a single source plane to good approximation. Taking this into consideration, the lensing kernel for galaxy lensing radially integrates over the source distribution, giving

$$W_{\kappa, \text{gal}}(\chi) = \int_\chi^\infty d\chi' n(z') H(z') W_\kappa(\chi, \chi'), \quad (2.11)$$

where the distribution of source galaxies, n , is given as function of redshift, z , and is thus accompanied by the Hubble parameter, $H(z)$, as the conversion factor to χ .

Throughout this paper, we will only be using the lensing convergence within the Born approximation. The leading post-Born corrections to the convergence power spectrum are at the 0.2% level [11, 12] and are negligible for our analysis.

The accuracy of the convergence power spectrum is significantly improved at large scales using the extended Limber approximation, $k \approx (L + 1/2)/\chi$ [22]. However, we use the lowest order Limber approximation for modelling of second order post-Born statistics, such as rotation (which we will describe in terms of the convergence power spectrum), hence for consistency we keep eq. (2.8) to lowest order.

2.2 Post-Born rotation

To get field rotation we need to go beyond the Born approximation and into the post-Born regime. Looking at the second order term in eq. (2.4) we see that the deflection field is no longer curl free, and thus cannot be fully described solely by the lensing potential. We introduce the pseudo-scalar *curl potential* Ω [14] so that

$$\boldsymbol{\alpha}(\hat{\mathbf{n}}) = \nabla \phi(\hat{\mathbf{n}}) + \nabla \times \Omega(\hat{\mathbf{n}}). \quad (2.12)$$

Analogous to the lensing convergence, the curl field is accessible through the deflection field via the *lensing rotation* observable

$$\omega(\hat{\mathbf{n}}) = -\frac{1}{2} \nabla \times \boldsymbol{\alpha}(\hat{\mathbf{n}}) = -\frac{1}{2} \nabla^2 \Omega(\hat{\mathbf{n}}). \quad (2.13)$$

This describes the rotational distortion to a field, and first appears at second order due to the *lens-lens* interaction of a particular photon bundle. That is, the result of two sequential lensing events (specifically, two unaligned shearings) induces rotation. The explicit form of field rotation in harmonic space (again in the flat sky limit) is

$$\omega(\mathbf{L}) = -2 \int_0^{\chi_s} d\chi W(\chi, \chi_s) \int_0^\chi d\chi' W(\chi', \chi) \times \int \frac{d^2 \mathbf{l}}{(2\pi)^2} (\mathbf{l} \cdot \mathbf{l}') [\mathbf{l} \times \mathbf{L}] \Psi(\mathbf{l}, \chi) \Psi(\mathbf{l}', \chi'), \quad (2.14)$$

where $\mathbf{l}' \equiv \mathbf{L} - \mathbf{l}$ for brevity.

Another second order effect comes from the *ray-lens* interaction. This correction is not due to a second lensing event, but instead to a change in the gravitational potential gradients along the deflected path of the photon. However, it produces no rotation, and we are not considering second order corrections to the convergence here.

³For the CMB lensing convergence, $\chi_s = \chi_*$. However we leave the kernel in this more general form as it is useful for the definition of the curl power spectrum eq. (2.15), and the galaxy lensing window function eq. (2.11).

2.2.1 Rotation power spectrum

The rotation power spectrum is [12]

$$C_L^{\omega\omega}(\chi_s) = \int \frac{d^2\mathbf{l}}{(2\pi)^2} \frac{(\mathbf{l} \cdot \mathbf{l}')^2 [\mathbf{l} \times \mathbf{L}]^2}{l^4 (l')^4} \mathcal{M}_{\chi_s}^{\omega\omega}(\mathbf{l}, \mathbf{l}'), \quad (2.15)$$

where the rotation mode-coupling matrix is defined as

$$\mathcal{M}_{\chi_s}^{\omega\omega}(\mathbf{l}, \mathbf{l}') = \int_0^{\chi_s} d\chi \frac{W_\kappa(\chi, \chi_s)^2}{\chi^2} P_{\delta\delta} \left(k \approx \frac{l}{\chi}, z(\chi) \right) C_{l'}^{\kappa\kappa}(\chi, \chi). \quad (2.16)$$

We only consider the auto spectrum of lensing rotation estimators at the same redshift, as appropriate for the CMB, so $C_L^{\omega\omega}$ only depends on a single source plane.

For a more detailed description of the post-Born derivation and in-depth discussion of its impact, see ref. [12].

3 Is rotation detectable?

The rotation signal from post-Born lensing is a second order lensing effect, so it is subdominant to convergence and shear from the gradient deflection field. We now explore whether the rotation is actually detectable within the near future using current lensing estimator technology.

3.1 CMB lensing reconstruction

A fixed weak lensing field breaks the statistical isotropy of the CMB sky, introducing correlations between unequal modes [1]. These new correlations can be probed to extract information about the lensing field that produced them. This is typically done with quadratic estimators [23, 24] for their computational efficiency. In ref. [25], a global minimum variance (GMV) quadratic estimator (QE) for the lensing potential is constructed using quadratic combinations of CMB fields weighted by the inverse of their covariance. For lensed CMB fields $\tilde{X} \in \{\tilde{T}, \tilde{E}, \tilde{B}\}$, one can similarly construct a GMV estimator for the field rotation

$$\hat{\omega}(\mathbf{L}) = \frac{L^2}{2} A_L^\Omega \int \frac{d^2\mathbf{l}}{(2\pi)^2} \tilde{X}^i(\mathbf{l}) g_{ij}^\Omega(\mathbf{l}, \mathbf{L}) \tilde{X}^j(\mathbf{l}'), \quad (3.1)$$

with corresponding inverse covariance weighted functions

$$g_{ij}^\Omega(\mathbf{l}, \mathbf{L}) = \frac{1}{2} (\tilde{\mathbf{C}}_{\text{CMB}}^{-1})_l^{ip} f_{pq}^\Omega(\mathbf{l}, \mathbf{L}) (\tilde{\mathbf{C}}_{\text{CMB}}^{-1})_{l'}^{jq}. \quad (3.2)$$

The response functions, $f_{ij}^\Omega(\mathbf{l}, \mathbf{L})$, quantify the correlations between modes \mathbf{l} and \mathbf{l}' that are induced by the curl potential,

$$\left\langle \frac{\delta}{\delta\Omega(\mathbf{L})} \tilde{X}(\mathbf{l}) \tilde{Y}(\mathbf{l}') \right\rangle = \delta(\mathbf{l} + \mathbf{l}' - \mathbf{L}) f_{XY}^\Omega(\mathbf{l}, \mathbf{L}), \quad (3.3)$$

and are explicitly stated in table 1. The covariance matrices include the lensed CMB spectra plus noise, $(\tilde{\mathbf{C}}_{\text{CMB}})_{l'}^{ij} = \tilde{C}_l^{ij} + N_l^{ij}$. For simplicity we only consider complete flat-sky observations with isotropic noise, so that the covariances are diagonal in \mathbf{l} .

XY	$f_{XY}^{\Omega}(\mathbf{l}, \mathbf{L})$
TT	$\mathbf{l} \times \mathbf{L} \left(\tilde{C}_l^{T\nabla T} - \tilde{C}_l^{T\nabla T} \right)$
EE	$\mathbf{l} \times \mathbf{L} \left(\tilde{C}_l^{E\nabla E} - \tilde{C}_l^{E\nabla E} \right) h_E(\mathbf{l}, \mathbf{l}')$
EB	$\mathbf{l} \times \mathbf{L} \left(\tilde{C}_l^{E\nabla E} - \tilde{C}_l^{B\nabla B} \right) h_B(\mathbf{l}, \mathbf{l}')$
TE	$\mathbf{l} \times \mathbf{L} \left(\tilde{C}_l^{T\nabla E} h_E(\mathbf{l}, \mathbf{l}') - \tilde{C}_l^{T\nabla E} \right)$
TB	$\mathbf{l} \times \mathbf{L} \tilde{C}_l^{T\nabla E} h_B(\mathbf{l}, \mathbf{l}')$

Table 1. The curl potential response functions for individual pairs of CMB maps X and Y required for lensing reconstruction. The geometric factors are defined as $h_E(\mathbf{l}_1, \mathbf{l}_2) = \cos(2(\theta_{l_1} - \theta_{l_2}))$, and $h_B(\mathbf{l}_1, \mathbf{l}_2) = \sin(2(\theta_{l_1} - \theta_{l_2}))$, and $\mathbf{l}' = \mathbf{L} - \mathbf{l}$. The non-perturbative response functions given here (slightly generalizing previous results [14, 23]) use the lensed gradient spectra defined in ref. [26]. The BB response function is excluded along with the curl-like terms, $C^{XP\perp}$, as they are subdominant relative to the other terms [27].

The quadratic estimator is normalized by

$$[A_L^{\Omega}]^{-1} = [N_0^{\Omega}(L)]^{-1} = \int \frac{d^2\mathbf{l}}{(2\pi)^2} g_{ij}^{\Omega}(\mathbf{l}, \mathbf{L}) f_{ij}^{\Omega}(\mathbf{l}, \mathbf{L}), \quad (3.4)$$

which is conveniently related to the lowest order reconstruction noise, $N_0^{\omega}(L) = L^4 N_0^{\Omega}(L)/4$, on the reconstructed field rotation. The GMV estimator for the CMB lensing convergence is very similar, only differing slightly in the response functions (which can be found in ref. [25]).

For low noise levels the quadratic estimator becomes suboptimal, effectively being limited by the lensing variance in the lensed B modes. Iterative reconstruction methods can maximize the likelihood (or posterior) more directly, giving substantially lower reconstruction noise for noise levels where the B modes are signal dominated [28, 29]. For measurement of the convergence power spectrum, gains on cosmological parameters from using iterative reconstruction are generally modest, because the well-resolved modes are dominated by cosmic variance of the lenses. For lensing rotation, where the signal is very small (and zero in the null-hypothesis of no rotation), the rotation reconstruction is noise dominated so the signal variance is negligible. Iterative reconstruction can therefore give larger improvements. For our purposes, we just need a forecast for the iterative lensing reconstruction noise, which can be calculated straightforwardly by iteration of the quadratic estimator formula using partially delensed fields: see appendix A for details.

3.1.1 CMB experimental setups

In this paper we consider different configurations of two near future CMB experiments, SO (at baseline and goal sensitivities), and CMB-S4. For modelling the effective CMB map level noise from all detectors, N_l^X , we use the harmonic-space Internal Linear Combination (ILC) model, described in [15], and include detector, beam, and foreground contaminants.⁴ We

⁴For SO, the noise curves are taken from https://github.com/simonsobs/so_noise_models/tree/master/LAT_comp_sep_noise/v3.1.0. These include baseline and goal sensitivities. For CMB-S4, the curves are taken from http://sns.ias.edu/~jch/S4_190604d_2LAT_Tpol_default_noisecurves.tgz.

enforce effective cuts of $30 \leq L \leq 5000$ for the polarization maps, and $30 \leq L \leq 3000$ for the temperature. This is applied to all experimental configurations through setting $N_l^{ij} = \infty$ outside the specified range. The fractional sky coverage of both experiments is approximated as $f_{\text{sky}} = 0.4$.

The noise on the GMV reconstructed observables are modelled at lowest order, i.e. with eq. (3.4) for $\hat{\omega}$, and a similar result for $\hat{\kappa}$ [25]. We also include forecasts with maximum a posteriori (iterative) estimators in our analyses, which have reconstruction noise calculated with the PLANCKLENS code [2]. In general, we do not assume external tracers are used in the delensing step of the iterative reconstruction.

We also include a more speculative experimental setup, which we label ‘CMB-S4 deep’. Consider a CMB-S4 measurement covering a small patch of the sky, say $f_{\text{sky}} = 0.05$. Long exposure to such a small area could reduce the polarization map level noise to $\Delta P = 0.5 \mu\text{K}$ -arcmin. If we assume this particular patch is free of any significant foreground contamination then we can neglect foreground cleaning and simply model the noise by [30]

$$N_L^P = \left(\frac{\Delta P \times 10^{-6}}{T_{\text{CMB}}} \right)^2 \exp \left(\frac{L(L+1)\theta^2}{8 \ln 2} \right), \quad (3.5)$$

which accounts for deconvolving a beam of full width half maximum (FWHM) $\theta = 2.3$ arcmin. It is harder to justify the exclusion of foreground modelling on temperature measurements, hence for this configuration we use polarization-only reconstruction. For consistency, we keep the multipole cuts of $30 \leq L \leq 5000$.

Figure 1 shows the polarization noise power spectra expected for the different experimental setups. The CMB-S4 deep can be considered a likely limiting best case scenario for the CMB-S4 lensing reconstruction noise. It is the only configuration with signal-dominated C_L^{BB} at any scale, and is therefore ideal for iterative reconstruction methods. The large-patch foreground-cleaned configurations, by contrast, provide conservative results due to the substantial temperature and polarization foreground power assumed over the entire area.

3.2 CMB rotation forecasts

It is instructive to consider the theoretical limit on the rotation signal-to-noise ratio, S/N , for a given experiment. We do this by constructing forecasts via the Fisher information [31, 32]. Using curl reconstruction from observed lensed CMB maps, we assess detectability by how well we can measure the amplitude of the reconstructed rotation power spectrum, $C_L^{\hat{\omega}\hat{\omega}}$, compared to the fiducial model. Approximating the reconstruction noise as Gaussian and isotropic, and working in the flat-sky approximation, the Fisher estimate reduces to [33]

$$\left(\frac{S}{N} \right)_{\hat{\omega}\hat{\omega}}^2 \leq f_{\text{sky}} \int L dL \frac{(C_L^{\omega\omega})^2}{(C_L^{\omega\omega} + N_L^\omega)^2}. \quad (3.6)$$

Here we only include reconstruction noise at lowest order $N_L^\omega = N_0^\omega(L)$. Inclusion of the next leading noise term, N_1^ω , could be important on certain scales, and degrade the forecasts further. However, it is still expected to be subdominant to N_0^ω and is not considered here. The integral is evaluated over the multipole range $30 \leq L \leq 4000$.

The significance at which the reconstructed CMB lensing rotation power spectrum could be detected is shown in the first column of table 2 for the GMV estimator,⁵ and iterative

⁵Note, as the CMB-S4 deep configuration only uses polarization maps, the quadratic estimation is technically not a GMV reconstruction. However, the methods are identical, just without input from the temperature.




Figure 1. The assumed effective CMB polarization noise power on E-mode (solid lines) and B-mode (dotted lines) measurements for different SO- and CMB-S4-like experimental setups. The lensed E-mode power-spectrum (dot-dashed line) is signal dominated on large scales for all configurations considered. However, the lensed B-mode spectrum (dashed line) is noise dominated on all scales with the wide area configurations. CMB-S4 deep patch assumes noise levels significantly below the B-mode power on large scales, but only covers 5% of the sky. The deep patch assumes the limiting best case of white noise with no increase in noise due to foreground cleaning.

Experiment	$C^{\hat{\omega}\hat{\omega}} [\sigma]$	$C^{\hat{\omega}\omega} [\sigma]$	$C^{\hat{\omega}^{\text{gal}}\hat{\omega}^{\text{gal}}} [\sigma]$	$C^{\hat{\omega}^{\text{gal}}\omega^{\text{gal}}} [\sigma]$
SO baseline	0.0 (0.0)	9.6 (9.8)	—	—
SO goal	0.1 (0.1)	11.7 (12.0)	—	—
CMB-S4	0.4 (0.4)	23.5 (24.7)	—	—
CMB-S4 deep	0.5 (1.7)	17.9 (28.4)	—	—
LSST	—	—	0.0	3.3
Far future	—	—	0.0	8.3

Table 2. Fisher forecasts for the expected detection significance (signal-to-noise ratios, S/N) of different field rotation power spectra with near future experiments/surveys. The first column is for the auto-spectrum of the CMB reconstructed field rotation, $\hat{\omega}$, and the second column shows the potential improvement achievable by crossing it with a perfect tracer for the rotation field. The S/N numbers in brackets are the result of using iterative estimators instead of the GMV quadratic estimator. The CMB-S4 deep forecasts assume polarization-only reconstruction, and a fractional sky coverage of $f_{\text{sky}} = 0.05$. The last two columns show the corresponding forecasts for observations of galaxy lensing field rotation, $\hat{\omega}^{\text{gal}}$, assuming a single broad redshift bin.

reconstruction in brackets. It is clear the reconstructed rotation auto-spectrum will be very challenging to detect with the next generation of CMB experiments, in agreement with ref. [12]: the detection significance for SO at goal sensitivity is $\leq 0.1\sigma$, and there is only




Figure 2. The lowest-order reconstruction noise power, N_0^ω , for GMV curl reconstruction (dashed lines), and iterative reconstruction (dotted lines) with SO- and CMB-S4-like experimental configurations. The CMB-S4 deep setup assumes polarization-only reconstruction, and has 5% sky coverage. The noise levels are typically ~ 3 –4 orders of magnitude above the fiducial post-Born rotation power spectrum signal expected (solid black). The weak lensing shape noise for an LSST-like survey (red dot dashed line) is at least 5 orders of magnitude greater than the expected post-born rotation signal from galaxy lensing (solid grey).

marginal improvement for CMB-S4 with predicted best S/N at 0.4σ or 0.5σ with deep patch (ref. [12] forecast 0.7σ for S4, however their multipole cuts were different, and they use a simpler noise model without foreground cleaning). Iterative reconstruction does provide significant improvement for the CMB-S4 deep patch configuration, in which the auto-spectrum is forecast to be (un)detectable at 1.7σ . Here the B-mode lensing signal benefits from the low polarization noise, giving rise to efficient delensing during the iterative process.

The difficulty of direct rotation detection is illustrated in figure 2: the reconstruction noise from the quadratic estimators is several orders of magnitude larger than the post-Born rotation power spectrum for all future CMB experiments considered.

Now consider the cross correlation of the lensing reconstructed rotation, $\hat{\omega}$, with some external rotation template, $\hat{\omega}^{\text{ext}}$. In the limit that this external rotation is a perfect template of the truth, i.e. $\hat{\omega}^{\text{ext}} = \omega$, then the best case $(S/N)^2$ is now

$$\left(\frac{S}{N}\right)_{\hat{\omega}\omega}^2 \leq 2f_{\text{sky}} \int LdL \frac{C_L^{\omega\omega}}{2C_L^{\omega\omega} + N_L^\omega}. \quad (3.7)$$

The results in table 2 show this cross-correlation signal to be easily detectable with a sufficiently accurate external template for field rotation. In the limit that a perfect rotation template can be constructed, a significant detection of 12σ is possible for SO, and improves by a factor of 2 for CMB-S4. Iterative reconstruction is of marginal importance at SO noise levels, however can improve the detectability by $> 1\sigma$ for CMB-S4. Again, the deep patch CMB-S4

benefits significantly from iterative reconstruction, and can produce a cross-spectrum signal at 28.4σ , $\sim 4\sigma$ greater than the wide CMB-S4 forecast.

The challenge to the detection of post-Born rotation on the CMB is now in the construction of a sufficiently accurate rotation template. The creation of such a template using density tracers is explored in section 4.

3.3 Galaxy lensing rotation

Post-born lensing also contributes a non-vanishing curl component to galaxy lensing maps. However, galaxy lensing is mainly restricted to $z \lesssim 3$ as source galaxy numbers are only significant at low redshifts. It is therefore expected that post-Born corrections will be less important for galaxy lensing due to the reduced path length, and the rotation signal specifically will be weakened [11]. Here we make a brief argument for focussing on lensing rotation with the CMB over galaxy lensing.

Assuming a standard galaxy shear estimator is used to measure shear B modes (equivalent to rotation), the main source of variance is from shape noise [34]

$$N_L^{\omega^{\text{gal}}} \sim N_{\text{shape}}^s = \frac{\sigma_\epsilon}{\bar{n}}, \quad (3.8)$$

where $\sigma_\epsilon = 0.21$ is the standard deviation of galaxy ellipticities per component [35], and \bar{n} is the mean galaxy density per steradian. We consider two future galaxy surveys: the LSST gold sample with $\bar{n} = 40$ galaxies per arcminute [36] and $f_{\text{sky}} = 0.4$, and some “far future” full sky survey with $\bar{n} = 100$ galaxies per arcminute. Both models share the same radial galaxy distribution described in appendix B by eq. (B.2) and eq. (B.3). Then eq. (3.6) and eq. (3.7) produce the galaxy lensing rotation forecasts shown in the last two columns of table 2.

As expected, the detection of the rotation auto-spectrum from galaxy surveys is even more challenging than for CMB experiments. For a near future LSST-like survey, a statistically significant measurement will be difficult even with a perfect rotation tracer. Realistic improvements to detectability require substantial reductions to the lensing noise, as the signal is more than 5 orders of magnitude smaller than current noise levels, see figure 2. Near full sky coverage will also boost the signal as demonstrated with the “far future” configuration, however even that is only forecast to have an upper limit of 8.3σ with a perfect external rotation tracer.

These simplistic forecasts only model the lensing signal from one effective redshift bin. Improvements could be gained with more realistic modelling of the curl signal from the full tomography expected for upcoming surveys. Other methods using additional polarization measurements for each galaxy source where available, such as those described in ref. [37], could also improve galaxy rotation constraints. We leave these considerations for future work. For the remainder of this paper we focus on detecting field rotation from CMB lensing only, where the signal is larger and a detection is more easily achievable.

4 Optimal rotation template

In section 3.2, we showed that while the field rotation auto-spectrum will be undetectable for near-future CMB experiments, the correlation between CMB lensing and a suitable rotation template is detectable. Post-Born rotation is generated at lowest order by two lensing events (see section 2.2). That is, image rotation is produced by the combination of a linear lensing shear followed by a subsequent (non-aligned) lensing shear at a lower redshift. The

linear shear field is directly related to the convergence field, since they are both obtained as derivatives of an underlying scalar lensing potential, and the local source for the convergence depends on the matter density. This suggests that by combining tracers of the density that probe at least two different redshifts, we should be able to construct a template for the expected rotation signal. The CMB lensing convergence has significant contributions at all redshifts, so two copies of the convergence reconstruction could work. External density tracers are generally limited to lower redshifts ($z < 3$), but can have much higher signal-to-noise and better redshift resolution. To extract the most information, we allow for general combinations of the convergence and external tracers.

4.1 Template construction

Since the post-Born rotation template is expected to be quadratic in the density tracers, we can use a quadratic estimator in the form of eq. (3.1). For a set of tracers $\{\hat{a}^i\}$, we define a general weighted quadratic estimator

$$\hat{\omega}^{\text{tem}}(\mathbf{L}) \equiv \frac{F_L^{-1}}{2} \int \frac{d^2 \mathbf{L}_1}{(2\pi)^2} G_{ij}(\mathbf{L}_1, \mathbf{L}) \hat{a}^i(\mathbf{L}_1) \hat{a}^j(\mathbf{L} - \mathbf{L}_1), \quad (4.1)$$

for some weights G_{ij} . The normalization condition is set such that the cross-correlation with the true ω returns the fiducial post-Born rotation spectrum

$$\langle \hat{\omega}^{\text{tem}}(\mathbf{l}) \omega(\mathbf{l}') \rangle \equiv (2\pi)^2 \delta_D(\mathbf{l} + \mathbf{l}') C_l^{\omega\omega}. \quad (4.2)$$

Using this condition, and minimizing the variance to find the optimal weight functions, gives a normalization

$$F_L = \frac{1}{2C_L^{\omega\omega}} \int \frac{d^2 \mathbf{L}_1}{(2\pi)^2} b_{(-\mathbf{L})\mathbf{L}_1\mathbf{L}_2}^{\omega ij} (\mathbf{C}_{\text{LSS}}^{-1})_{\mathbf{L}_1}^{ip} (\mathbf{C}_{\text{LSS}}^{-1})_{\mathbf{L}_2}^{jq} b_{(-\mathbf{L})\mathbf{L}_1\mathbf{L}_2}^{\omega pq}, \quad (4.3)$$

with the minimum variance template estimator given explicitly by

$$\hat{\omega}^{\text{tem}}(\mathbf{L}) = \frac{F_L^{-1}}{2} \int \frac{d^2 \mathbf{L}_1}{(2\pi)^2} b_{(-\mathbf{L})\mathbf{L}_1\mathbf{L}_2}^{\omega ij} \hat{a}^i(\mathbf{L}_1) \hat{a}^j(\mathbf{L}_2). \quad (4.4)$$

Here we defined inverse covariance weighted fields $\bar{\mathbf{a}} \equiv (\mathbf{C}_{\text{LSS}}^{-1}) \mathbf{a}$, with the covariance having signal and noise contributions $(\mathbf{C}_{\text{LSS}})^{ij}_L = C_L^{ij} + N_L^{ij}$. The optimal weights are given in terms of the flat-sky rotation bispectra [38] defined by

$$\langle \omega(\mathbf{l}) a_i(\mathbf{l}_1) a_j(\mathbf{l}_2) \rangle = (2\pi)^2 \delta(\mathbf{l} + \mathbf{l}_1 + \mathbf{l}_2) b_{\mathbf{l}\mathbf{l}_1\mathbf{l}_2}^{\omega ij}. \quad (4.5)$$

As these bispectra have odd parity [12], the moduli of the modes are insufficient to determine the bispectrum sign. Hence, the sign of the modes in eq. (4.3) and eq. (4.4) are important to determine the handedness of the triangle and ensure $\mathbf{L}_2 = \mathbf{L} - \mathbf{L}_1$.

A useful consequence of the chosen normalization for the estimator in eq. (4.2) is that the variance of the estimator is

$$C_L^{\hat{\omega}^{\text{tem}} \hat{\omega}^{\text{tem}}} = C_L^{\omega\omega} + N_L^{\hat{\omega}^{\text{tem}}} = C_L^{\omega\omega} F_L^{-1}. \quad (4.6)$$

This in turn provides a relation between the normalization, F_L , and the correlation coefficient of the estimator with the true rotation

$$\rho_{\omega \hat{\omega}^{\text{tem}}}^2(L) = F_L. \quad (4.7)$$

4.2 Tracers

We assume that density tracers used as input to the template estimator are Gaussian fields with uncorrelated noise. If we restrict the form of these tracers to

$$a(\mathbf{L}) = \int d\chi W_a(\chi) \delta(\mathbf{L}, \chi), \quad (4.8)$$

then they are fully described by their window function, W_a , which describe the radial sensitivity to the underlying matter distribution. The general form of the density auto- and cross-spectra at lowest order Limber (for $a \neq \kappa$) is then

$$C_L^{a_1 a_2} = \int d\chi' \frac{W_{a_1}(\chi') W_{a_2}(\chi')}{(\chi')^2} P_{\delta\delta} \left(k \approx \frac{L}{\chi'}, z(\chi') \right), \quad (4.9)$$

or when the cross-spectrum involves one κ

$$C_L^{a\kappa}(\chi) = \int d\chi' \frac{W_a(\chi') W_\kappa(\chi', \chi)}{(\chi')^2} P_{\delta\delta} \left(k \approx \frac{L}{\chi'}, z(\chi') \right). \quad (4.10)$$

For completeness, the CMB convergence auto-spectrum is given in eq. (2.8).

The rotation-density-density bispectra in eq. (4.4) can be written explicitly as

$$b_{(-\mathbf{L})\mathbf{L}_1\mathbf{L}_2}^{\omega ij} = 2 \frac{(\mathbf{L}_1 \cdot \mathbf{L}_2) [\mathbf{L}_1 \times \mathbf{L}_2]}{(L_1)^2 (L_2)^2} [\mathcal{M}^{ij}(L_1, L_2) - \mathcal{M}^{ji}(L_2, L_1)], \quad (4.11)$$

with $\mathbf{L}_2 = \mathbf{L} - \mathbf{L}_1$, where the mode couplings are defined as

$$\mathcal{M}^{a_1 a_2}(L, L') \equiv \int_0^{\chi_*} d\chi \frac{W_\kappa(\chi, \chi_*) W_{a_1}(\chi)}{\chi^2} P_{\delta\delta} \left(k \approx \frac{L}{\chi}, z(\chi) \right) C_{L'}^{a_2 \kappa}(\chi). \quad (4.12)$$

As we are now only considering CMB lensing, it is implicitly assumed that $\chi_s = \chi_*$. Also, in the case of $a_2 = \kappa$ in eq. (4.12), $C_{L'}^{a_2 \kappa}(\chi) = C_{L'}^{\kappa \kappa}(\chi, \chi_*)$ which is given in eq. (2.8). The form of these bispectra are in general agreement to the similarly-defined post-Born convergence bispectra with two external tracers [39].

These rotation bispectra should be individually more detectable than the rotation reconstruction auto-spectrum, as good density tracers will not suffer from the severe noise domination that afflicts the lensing-reconstructed $\hat{\omega}$. This has already been demonstrated for the rotation-convergence-convergence bispectrum in ref. [12].

4.2.1 Tracer choices

We consider three density tracers as input for the rotation template: the CMB lensing convergence, κ , galaxy clustering, g , and the CIB, I . The modelling of these tracers, specifically of their window functions, along with their noise contributions are detailed in appendix B. Combinations of tracers that probe different redshift regions are ideal for a post-Born rotation template, as they are better able to mimic the lens-lens coupling. From the window functions plotted in figure 3, it is clear the galaxy clustering observable probes the low redshift universe, only sensitive to the underlying matter field for $z \lesssim 3$. Here we consider unbinned g , however tomographic binning of the galaxy counts would in principle provide a better probe of the contributions to the rotation signal from low redshifts. The CIB has sensitivity peaked at higher redshift. As the last scattering surface is in the early universe, CMB fields are lensed by density perturbations over almost the entire redshift range, as seen in the broad window




Figure 3. The Window functions, W , for tracers of the underlying matter distribution: CMB lensing convergence (κ), galaxy counts (g), and CIB (I). All are functions of comoving radial distance, χ , with maxima normalized to unity, and have been converted into kernels for the Weyl gravitational potential via Poisson factor, γ , defined in eq. (2.9). In matter domination the Weyl potential is constant (before decaying slightly once dark energy is important), so the figure approximately illustrates the relative contributions of the different redshift regions traced by each observable.

function. The lensing kernels determining the lens-lens coupling are broad in redshift, so combinations of these tracers should provide reasonable coverage of the pairs of redshifts that contribute to the rotation signal. We do not consider cosmic shear as one of the input tracers here, though in principle it could also be used as a lower-redshift tracer.

For $\hat{\kappa}$, we continue to use the expected experimental configurations for SO and CMB-S4. The LSST gold model [36] is the natural choice for \hat{g} , as the survey sky coverage will optimally overlap with SO and both surveys are expected to be operational at around the same time period. Finally, \hat{I} is modelled from the currently-available Planck CIB 353 GHz map. Figure 4 shows the correlation coefficients of the different tracers with the true CMB lensing convergence, given the assumed experimental configuration

$$\rho_{\hat{a}\kappa}(L) = \frac{C_L^{a\kappa}}{\sqrt{(C_L^{aa} + N_L^a)C_L^{\kappa\kappa}}}. \quad (4.13)$$

The SO $\hat{\kappa}$ becomes noise dominated at high L , resulting in the poor correlation on small scales. This is improved by CMB-S4, however the trend is still visible. It is also seen that iterative reconstruction is most important for CMB-S4 deep patch, due to the expected dominance of the B-mode lensing signal. Galaxy clustering is signal dominated on all scales due to the high expected LSST galaxy number density, reflected in the consistently strong correlation with κ . The CIB has better correlation on small scales than $\hat{\kappa}$ from the wide




Figure 4. The correlation coefficient, ρ , of the fiducial CMB lensing convergence, κ , with LSS density observables: LSST gold sample galaxy counts, \hat{g} , and Planck 353 GHz CIB, \hat{I} . The CIB observable is shown with multipole cuts at $110 \leq L \leq 2000$, illustrated with the dashed lines. The coefficient between κ and the GMV reconstructed CMB lensing convergence, $\hat{\kappa}$, is shown for SO at baseline sensitivity and CMB-S4. The improved correlation obtained using iteratively-reconstructed $\hat{\kappa}$ is shown in the dotted lines. The CMB-S4 deep configuration assumes polarization-only reconstruction.

patch CMB experiments, however we apply strict modal cuts to the CIB of $110 \leq L \leq 2000$ due to uncertainty in modelling the dust and instrumental noise, see appendix B for more details. The correlation coefficients are consistent with refs. [40, 41].

4.3 Fast real-space template construction

The rotation template of eq. (4.4) involves a computationally expensive double integral over a pair of weighted external density maps. The estimator’s response functions are rotation bispectra, $b^{\omega ij}$, which from eq. (4.11) and eq. (4.12) involve radial integration over terms separable in \mathbf{L}_1 and \mathbf{L}_2 . Therefore, at each step along this integral, the template actually takes the form of a (weighted and normalized) Fourier-space convolution. We exploit this by using convolution theorem to write the integrand of the template as configuration-space products of mode-separated functions. Denoting a Fourier transform with, \mathcal{F} , the template can be quickly computed in the form

$$\hat{\omega}^{\text{tem}}(\mathbf{L}) = 2F_L^{-1} \int_0^{\chi_*} d\chi \frac{W_\kappa(\chi, \chi_*)}{\chi^2} \delta_{pq} \epsilon_{rs} \mathcal{F}[\Gamma^{pr} \Lambda^{qs}](\mathbf{n}, \chi), \quad (4.14)$$

where the functions $\Gamma = \mathcal{F}^{-1}(\gamma)$, and $\Lambda = \mathcal{F}^{-1}(\lambda)$ are defined as

$$\gamma^{pr}(\mathbf{L}, \chi) \equiv \frac{L^p L^r}{L^2} W_i(\chi) P_{\delta\delta} \left(k \approx \frac{L}{\chi}, z(\chi) \right) (\mathbf{C}_{\text{LSS}}^{-1})_L^{ij} a_j(\mathbf{L}), \quad (4.15)$$

$$\lambda^{pr}(\mathbf{L}, \chi) \equiv \frac{L^p L^r}{L^2} C_L^{i\kappa}(\chi) (\mathbf{C}_{\text{LSS}}^{-1})_L^{ij} a_j(\mathbf{L}). \quad (4.16)$$

Experiment	$b^{\hat{\omega}\hat{\kappa}\hat{\kappa}}$ [σ]	$b^{\hat{\omega}\hat{g}\hat{g}}$ [σ]	$b^{\hat{\omega}\hat{g}\hat{\kappa}}$ [σ]	$b^{\hat{\omega}\hat{I}\hat{I}}$ [σ]	$b^{\hat{\omega}\hat{I}\hat{\kappa}}$ [σ]	$b^{\hat{\omega}\hat{I}\hat{g}}$ [σ]
SO baseline	0.4 (0.5)	0.2 (0.3)	2.8 (2.9)	0.1 (0.1)	0.4 (0.5)	3.5 (3.5)
SO goal	0.6 (0.7)	0.2 (0.3)	3.8 (3.9)	0.1 (0.1)	0.6 (0.6)	4.3 (4.4)
CMB-S4	1.8 (2.0)	0.6 (0.6)	9.5 (10.3)	0.2 (0.2)	1.4 (1.6)	8.8 (9.3)
CMB-S4 deep	1.5 (3.2)	0.5 (0.7)	8.2 (14.8)	0.2 (0.3)	1.1 (2.1)	6.5 (11.0)

Table 3. Fisher forecasts for the detection significance (signal-to-noise ratios, S/N) of various mixed lensing rotation-tracer-tracer bispectra individually. The CMB field rotation, $\hat{\omega}$, and convergence, $\hat{\kappa}$, are reconstructed with the GMV estimator (or iterative method in brackets). The forecasts are given for noise levels expected of SO (at baseline and goal sensitivities) and CMB-S4. The additional CMB-S4 deep patch experiment uses polarization only reconstruction, and covers just 5% of the sky. The other observables are LSST gold sample galaxy counts, \hat{g} , and the Planck CIB measurements, \hat{I} , at 353 GHz.

Here L^i is component i of the vector \mathbf{L} . The template is now computationally tractable. Convergence is achievable with naive uniform sampling of ~ 100 steps in χ , which can be easily reduced with targeted sampling along the radial path.

5 Forecasts

5.1 Bispectra forecasts

Our aim is to probe the CMB post-Born rotation power spectrum, which we estimate through the cross-correlation of the CMB lensing reconstructed rotation, $\hat{\omega}$, with an externally constructed template, $\hat{\omega}^{\text{tem}}$. From eq. (4.4), it is apparent that $C_L^{\hat{\omega}\hat{\omega}^{\text{tem}}}$ is dependent on weighted combinations of observed bispectra $b_{\mathbf{L}\mathbf{L}_1\mathbf{L}_2}^{\hat{\omega}\hat{g}\hat{g}}$. It is therefore useful to understand the relative contributions that each rotation bispectrum provides toward the rotation power spectrum detection significance.

We consider forecasts of the S/N for the individual rotation bispectra via the flat sky Fisher information, where the parameter of interest is the amplitude of the bispectrum

$$\left(\frac{S}{N}\right)_{\hat{\omega}\hat{a}_1\hat{a}_2}^2 \leq \frac{f_{\text{sky}}}{\pi} \int \frac{d^2\mathbf{L}_1 d^2\mathbf{L}_2}{(2\pi)^2} \frac{(b_{\mathbf{L}\mathbf{L}_1\mathbf{L}_2}^{\omega a_1 a_2})^2}{\left[(\mathbf{C})_{L_1}^{a_1 a_1} (\mathbf{C})_{L_2}^{a_2 a_2} + (\mathbf{C})_{L_1}^{a_1 a_2} (\mathbf{C})_{L_2}^{a_1 a_2}\right] (C_L^{\omega\omega} + N_L^\omega)}. \quad (5.1)$$

As for the power spectra forecasts, the covariance matrix elements are $(\mathbf{C})_L^{a_1 a_2} = (\mathbf{C}_{\text{LSS}})_{L_1}^{a_1 a_2} = C_L^{a_1 a_2} + N_L^{a_1 a_2}$, and the integrals are evaluated over multipole range $30 \leq L \leq 4000$, with noise taken to be uncorrelated between observables, i.e. $N_L^{a_1 a_2} = 0$ if $a_1 \neq a_2$.

The resulting upper-bound S/N for the different bispectra are given in table 3. The pure CMB rotation bispectrum $b^{\omega\kappa\kappa}$ has an expected S/N of 3.2σ for iteratively reconstructed CMB-S4 deep patch, while the more conservative CMB-S4 wide result is only 2.0σ . Therefore, it will remain challenging for near-future CMB measurements alone to detect ω with any significance due to large reconstruction noise levels.

The most detectable bispectra combine a low redshift tracer with a high redshift one. It follows that bispectra involving galaxy density combined with either the CIB or κ , are thereby easier to detect. For SO baseline sensitivity, $b^{\hat{\omega}\hat{g}\hat{\kappa}}$ and $b^{\hat{\omega}\hat{g}\hat{I}}$ have signals at 2.8σ and 3.5σ significance. This is increased by $\sim 1\sigma$ if the goal sensitivity is reached. For CMB-S4, the




Figure 5. Contributions to the Fisher information as a function of the $\omega(\mathbf{L})$ rotation multipole (L) for selected bispectra from table 3. The curves are scaled by L to avoid contributions appearing misleading on smaller scales due to the choice of semi-logarithmic axes. The solid lines represent CMB reconstruction using GMV estimators, while the dotted show iterative reconstruction.

detection significance improves significantly to 10.3σ and 9.3σ with iterative reconstruction. This is further improved with a deep patch to 14.8σ and 11.0σ . Bispectra that combine observables occupying similar redshift space are unlikely to be detected even by CMB-S4. This is especially true for combinations of the same observable, but also for $b^{\hat{\omega}\hat{\kappa}}$ due to the significant overlap between the CIB and κ in redshift sensitivity as is shown in figure 3.

The Fisher contributions per log mode for each individual bispectra are shown in figure 5 for SO and CMB-S4 noise levels. The peak of the contributions to S/N for a pair of κ tracers is centred at $L \sim 500$, with very little contribution beyond $L > 2000$ due to lensing reconstruction noise dominance on small scales, as illustrated in figure 12. In contrast, the expected LSST galaxy density is signal dominated to high L , and hence bispectra involving g have significant contributions to S/N at $L > 1000$. For bispectra involving the CIB, contributions outside the stricter limits of $110 \leq L \leq 2000$ are suppressed.

5.2 Cross-spectra forecasts

We can now finally quantify the ability of the post-Born rotation template, $\hat{\omega}^{\text{tem}}$, to reproduce the underlying field rotation spectrum when crossed with the noise-dominated CMB lensed rotation estimator, $\hat{\omega}$. The flat-sky Fisher for this rotation cross spectrum is

$$\left(\frac{S}{N}\right)_{\hat{\omega}\hat{\omega}^{\text{tem}}}^2 \leq \frac{f_{\text{sky}}}{\pi} \int d^2\mathbf{L} \frac{C_L^{\omega\omega}}{(C_L^{\omega\omega} + N_L^\omega)F_L^{-1} + C_L^{\omega\omega}}. \quad (5.2)$$

In the limit of a perfectly reconstructed template, i.e. $F_L = 1$, we recover the Fisher information for the rotation cross-spectrum with a perfect tracer (eq. (3.7)) as expected. Hence, the

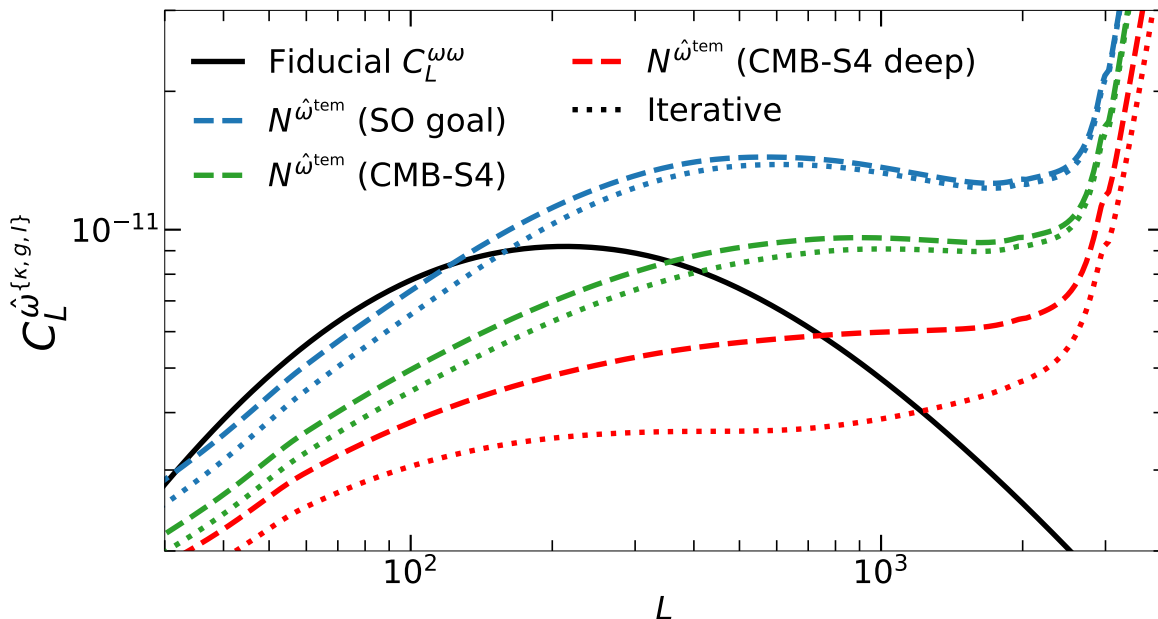


Figure 6. Noise levels (dashed lines) on the auto-spectrum of a rotation template, $\hat{\omega}^{\text{tem}}$, for SO at goal sensitivity, CMB-S4, and CMB-S4 deep patch. Dotted lines show corresponding results using iterative reconstruction on $\hat{\kappa}$ instead of GMV. Polarization-only reconstruction is used for CMB-S4 deep patch lensing observables, over a sky fraction of just 5%. The template is constructed from the optimal combination of all tracers considered in this paper, $\hat{a} \in \{\hat{\kappa}, \hat{g}, \hat{I}\}$. On large scales the template noise is low compared to the expected fiducial post-Born field rotation signal (solid black), but the template is noise-dominated on small scales.

limits originally set in section 3.2 hold true. However, if the template reconstruction is not perfect and the lensing reconstructed $\hat{\omega}$ is noise dominated, the Fisher information reduces to

$$\left(\frac{S}{N}\right)_{\hat{\omega}\hat{\omega}^{\text{tem}}}^2 \lesssim \frac{f_{\text{sky}}}{\pi} \int d^2\mathbf{L} \frac{C_L^\omega F_L}{N_L^\omega}. \quad (5.3)$$

This is the (flat sky equivalent) Fisher matrix of the standard optimal bispectrum estimator (see e.g. ref. [26]).

Does the template get close to a near-perfect reconstruction of the field rotation signal? Figure 6 plots the noise levels on the auto-spectrum of a template using all three density tracers, i.e. $\hat{a} \in \{\hat{\kappa}, \hat{g}, \hat{I}\}$, which we will denote $\hat{\omega}^{\{\kappa,g,I\}}$. It shows the noise levels are of similar magnitude to the estimated signal, a significant improvement compared to the noise-dominated CMB lensing reconstruction in figure 2. For CMB-S4 deep, the noise only begins to dominate at $L > 1000$ for a template using iteratively reconstructed $\hat{\kappa}$. However, noise does begin to dominate the template at lower L for the large patch experiments. Figure 7 shows the correlation coefficients of different templates with the fiducial rotation, $\rho_{\omega\hat{\omega}^{\text{tem}}}$, as defined in eq. (4.7). As expected, the best correlations are achieved when all the density tracers are used. The correlations are greatest for $L \lesssim 1000$, however F_L does not exceed 0.73 for any of the sensitivities considered. For near future experiments, $\hat{\omega}^{\text{tem}}$ will not achieve a close-to-perfect emulation at any mode, but should still recover significant portions of the rotation signal.

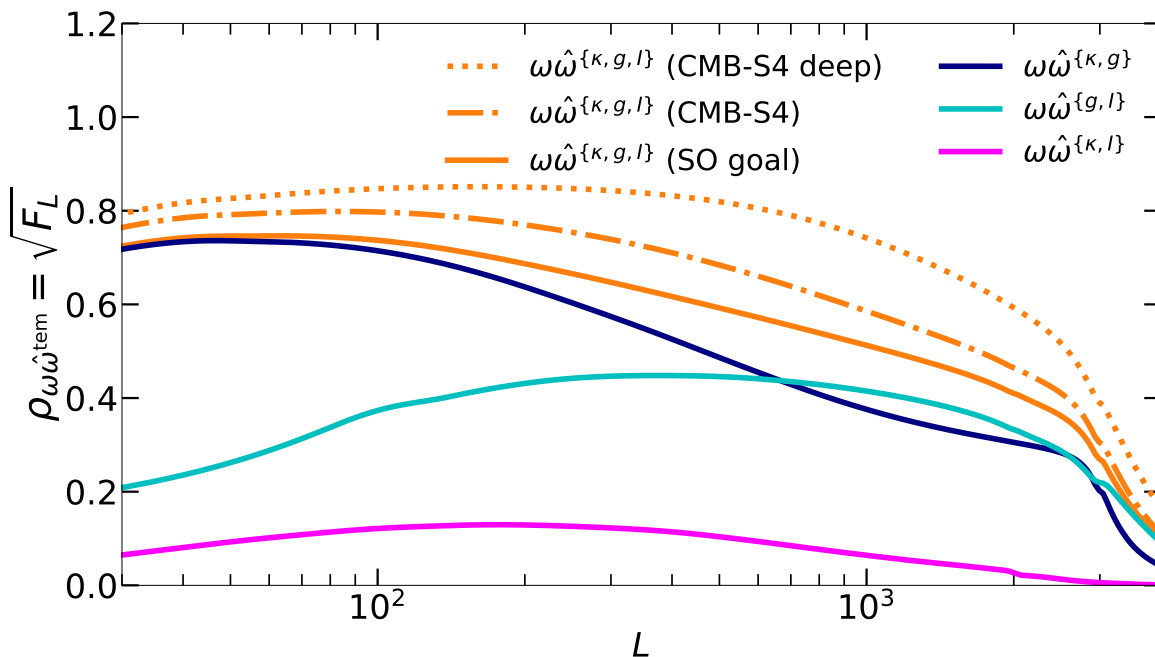


Figure 7. The correlation coefficients, ρ , for different rotation templates with the true fiducial post-Born field rotation, calculated using eq. (4.7). All results assume iterative reconstruction for the lensing convergence tracer, either with noise levels for SO goal (solid lines), CMB-S4 (dash-dotted line), or CMB-S4 deep (dotted line). CMB-S4 deep uses only polarization in the lensing reconstructions.

As lensing reconstruction is completely noise dominated, $C_L^{\omega\omega} \ll N_L^\omega$, the Fishers used to forecast the upper-bound S/N on the rotation cross-spectra, eq. (5.2) and eq. (5.3), are equivalent. The results can therefore also be interpreted as forecasts for optimal rotation-density-density bispectra, and are shown in table 4 for different sets of tracers. A statistically significant detection of $> 5\sigma$ is now feasible with SO. Using LSST \hat{g} , the Planck 353 GHz \hat{I} , and CMB reconstructed $\hat{\kappa}$, the forecast reconstruction of the rotation spectrum is 4.5σ – 5.7σ for SO baseline-goal sensitivities. This is improved by 0.2σ – 0.4σ with iterative lensing reconstruction. For the CMB-S4 wide configuration, the signal can be probed at 13.6σ – 14.7σ depending on the lensing reconstruction method used. This is improved significantly with the deep patch configuration, resulting in the best forecast of 21.3σ when iterative methods are used for the lensing observables.

When only considering a set of two unique tracers for template construction, the trends seen in the bispectra forecasts (and correlation coefficients) are also seen in these cross-correlation forecasts. Optimally constructed templates $\hat{\omega}^{\{g,l\}}$ and $\hat{\omega}^{\{\kappa,g\}}$ produce the most detectable cross-spectra, with $> 4.5\sigma$ achievable for SO at goal sensitivity. When looking ahead to CMB-S4, $\hat{\omega}^{\{\kappa,g\}}$ gets the greatest boost due to the lower reconstruction levels in both $\hat{\kappa}$ and $\hat{\omega}$, and consequently is detectable at 12.1σ – 13.3σ . However, $\hat{\omega}^{\{g,l\}}$ only benefits from the noise reduction in $\hat{\omega}$, and so is slightly less enhanced at 9.3σ – 9.8σ . This is also seen for CMB-S4 deep patch with iterative reconstruction, in which $\hat{\omega}^{\{\kappa,g\}}$ and $\hat{\omega}^{\{g,l\}}$ are forecast at 20.4σ and 11.6 respectively. The template $\hat{\omega}^{\{\kappa,l\}}$ is only detectable at 3.7σ in the best case for CMB-S4 deep patch with iterative reconstruction.

The variance of the cross-spectrum for a bin centred at L with width ΔL is given by

$$\sigma_L^2 = \frac{1}{2L\Delta L f_{\text{sky}}} C_L^{\omega\omega} N_L^\omega F_L^{-1}. \quad (5.4)$$

Experiment	$C^{\hat{\omega}\hat{\omega}^{\{\kappa,g\}}} [\sigma]$	$C^{\hat{\omega}\hat{\omega}^{\{\kappa,I\}}} [\sigma]$	$C^{\hat{\omega}\hat{\omega}^{\{g,I\}}} [\sigma]$	$C^{\hat{\omega}\hat{\omega}^{\{\kappa,g,I\}}} [\sigma]$
SO baseline	3.3 (3.5)	0.6 (0.6)	3.6 (3.7)	4.5 (4.7)
SO goal	4.5 (4.9)	0.8 (0.9)	4.4 (4.6)	5.7 (6.1)
CMB-S4	12.1 (13.3)	2.1 (2.4)	9.3 (9.8)	13.6 (14.7)
CMB-S4 deep	10.7 (20.4)	1.8 (3.7)	6.9 (11.6)	11.3 (21.3)

Table 4. Fisher forecasts for the detection significance (signal-to-noise ratios, S/N) of the cross-spectrum between the CMB reconstructed field rotation, $\hat{\omega}$, and different optimally-constructed rotation templates, $\hat{\omega}^{\text{tem}}$. The density tracers used as input to the template are labelled in the superscript, e.g. the template denoted $\hat{\omega}^{\{\kappa,g,I\}}$ is optimally constructed from all combinations of the CMB lensing convergence, $\hat{\kappa}$, galaxy clustering, \hat{g} , and CIB, \hat{I} . The results are given for different CMB experimental setups assuming GMV quadratic estimators (or iterative reconstruction in brackets). CMB-S4 deep uses polarization-only reconstruction, and assumes a fractional sky coverage $f_{\text{sky}} = 0.05$ with no foregrounds.

The 1σ error forecasts for the various experimental configurations are shown in figure 8, assuming the optimal template estimator, $\hat{\omega}^{\{\kappa,g,I\}}$. The spectrum is constrained away from zero for all noise levels in the multipole range $200 < L < 3000$. At large scales the lack of modes (cosmic variance) dominates as is shown by the relatively larger error bars. Separately, on the small scales the CMB reconstruction noise starts to significantly dominate over the signal at $L \sim 3000$, also seen in figure 6. While the amplitude of the template spectrum depends on the galaxy bias, the detectability forecasts are largely independent of the bias due to the high S/N of \hat{g} . As long as the bias can be measured empirically (e.g. using cross-correlation with lensing convergence), the rotation template amplitude can be normalized empirically. Uncertainties in modelling of the galaxy bias (and potentially redshift distribution) only become relevant if the template amplitude is used for cosmological inference; however, the parameters are likely to be well constrained by convergence-density spectra compared to the uncertainty levels that would be relevant for the rotation.

There are a few avenues for improving these rotation forecasts further. For example, other density tracers could be used as input to the rotation template. Cosmic shear is the obvious candidate; although there would be significant overlap of redshift sensitivity with \hat{g} , cosmic shear is a more direct probe of the lensing effect. Another approach could be to use the full tomography of galaxy surveys (using galaxy counts and/or cosmic shear), giving a closer match to the ideal template. However, we only considered one large effective bin for \hat{g} here, as we found improvements with binning to be modest. Finally, using external tracers to improve the delensing in the iterative reconstruction can reduce the reconstruction noise further. For CMB-S4 deep, including \hat{g} and \hat{I} in the iterative reconstruction potentially gives an additional $\sim 2\sigma$ detection significance. However, the improvement is less apparent for the wide experimental setups. Using external tracers in the delensing may also complicate the cross-spectrum bias, since $\hat{\omega}^{\text{tem}}$ is constructed from the same (or correlated) tracers.

6 Cross-spectrum bias

The rotation-template cross spectrum is in general a biased estimator of the post-Born rotation power spectrum

$$\langle C^{\hat{\omega}\hat{\omega}^{\text{tem}}} \rangle = C^{\omega\omega} + \text{bias}. \quad (6.1)$$




Figure 8. The reconstruction error forecast on the template-rotation cross-spectrum using eq. (5.4). Error bands are shown for 17 logarithmically separated bins in L , and interpolated to produce smooth contours. Different shades represent different experimental configurations, and the fiducial post-Born rotation is displayed in solid orange. The CMB-S4 deep setup use polarization-only reconstruction, and has a sky coverage of just 5%; the increase in cosmic variance with smaller sky area is compensated on all scales by reduced rotation reconstruction noise due the lower instrumental noise, with additional gains from lower convergence tracer reconstruction noise and the limiting assumption of no foregrounds on the deep patch.

The bias is non-zero even when there is no deflection, so for accurate measurement of lensing rotation the bias must be accounted for. Large biases could potentially degrade the Fisher forecasts of the previous section. Here we explore the leading Gaussian contributions to this cross signal bias, and show that even configurations with bias of similar amplitude to $C^{\omega\omega}$ do not affect the S/N .

6.1 Template-lensing bias

CMB field rotation is reconstructed from pairs of lensed CMB fields, $\tilde{X}, \tilde{Y} \in \{\tilde{T}, \tilde{E}, \tilde{B}\}$, as shown in eq. (3.1). Likewise, the rotation template eq. (4.4) involves pairs of LSS tracers, e.g. $\hat{a}, \hat{b} \in \{\hat{\kappa}, \hat{g}, \hat{l}\}$. Therefore, the full expectation of the cross-spectrum, i.e. the l.h.s. of eq. (6.1), is proportional to the trispectrum $\langle \tilde{X}\tilde{Y}\hat{a}\hat{b} \rangle$.

As described in section 2, \tilde{X} is related to the unlensed field, X , by the deflection field given in eq. (2.5). Consider the series expansion of the lensed field at some point in sky in terms of the deflection field (ignoring curl terms here as they are subdominant to the lensing potential)

$$\begin{aligned} \tilde{X} &= X + \delta_\phi X + \delta_\phi^2 X + \dots \\ &= X + \nabla^i \phi \nabla_i X + \frac{1}{2} \nabla^i \phi \nabla^j \phi \nabla_i \nabla_j X + \dots \end{aligned} \quad (6.2)$$

Therefore, from perturbative expansion of the trispectrum, the leading⁶ terms present in the absence of rotation are

$$\langle \tilde{X}\tilde{Y}\hat{a}\hat{b} \rangle_{\text{bias}} = \langle \delta_\phi^2 XY \hat{a}\hat{b} \rangle + \langle X\delta_\phi^2 Y \hat{a}\hat{b} \rangle + \langle \delta_\phi X \delta_\phi Y \hat{a}\hat{b} \rangle + \mathcal{O}(\Psi^6). \quad (6.3)$$

These three terms contribute to the bias at the $N^{(2)}$ level⁷ and are similar to the $N^{(3/2)}$ bias terms in ref. [39] which arise from a non-negligible convergence bispectra in the CMB lensing-LSS cross correlation. However, here we are still simply assuming Gaussianity of all density tracers (including κ), so this bias is not caused by a LSS/lensing bispectrum. We can see that each term is actually proportional to a density trispectrum $\langle \phi\phi\hat{a}\hat{b} \rangle$, which we assume is dominated by the disconnected part. Alternatively, as we are using pairs of tracers to estimate the lensing rotation, so the rotation template is non-Gaussian, the bias could be also be considered as coming from template-lensing bispectra $\langle \phi\phi\hat{\omega}^{\text{tem}} \rangle$. We choose to follow the notational conventions of ref. [42] and label the first two terms of eq. (6.3) as $N_{\text{C1}}^{(2)}$, and the last as $N_{\text{A1}}^{(2)}$.

Assuming the unlensed CMB and late-time density tracers are independent (and Gaussianity of all the fields), using the perturbative expansion of the lensed CMB fields gives the analytic results

$$N_{\text{A1}}^{(2)}(L) = -\frac{L^2}{2} A_L^\Omega \int \frac{d^2\mathbf{l}d^2\mathbf{L}_1}{(2\pi)^4} b_{(-L)\mathbf{L}_1\mathbf{L}_2}^{\hat{\omega}^{\text{tem}}\phi\phi}(\mathbf{l}'' \cdot \mathbf{L}_1)(\mathbf{l}'' \cdot \mathbf{L}_2) \sum_{X,Y} g_{XY}^\Omega(\mathbf{l}, \mathbf{L}) C_{i''}^{\bar{X}\bar{Y}} h_X(\mathbf{l}'', \mathbf{l}) h_Y(\mathbf{l}'', \mathbf{l}'), \quad (6.4)$$

and

$$\begin{aligned} N_{\text{C1}}^{(2)}(L) &= \frac{L^2}{4} A_L^\Omega \int \frac{d^2\mathbf{l}d^2\mathbf{L}_1}{(2\pi)^4} b_{(-L)\mathbf{L}_1\mathbf{L}_2}^{\hat{\omega}^{\text{tem}}\phi\phi}(\mathbf{l} \cdot \mathbf{L}_1)(\mathbf{l} \cdot \mathbf{L}_2) \\ &\quad \times \sum_{X,Y} \left[g_{XY}^\Omega(\mathbf{l}, \mathbf{L}) C_i^{\bar{X}\bar{Y}} h_Y(\mathbf{l}, \mathbf{l}') + g_{YX}^\Omega(\mathbf{l}, \mathbf{L}) C_i^{\bar{X}\bar{Y}} h_X(\mathbf{l}, \mathbf{l}') \right], \end{aligned} \quad (6.5)$$

for $\mathbf{L}_2 = \mathbf{L} - \mathbf{L}_1$, $\mathbf{l}' = \mathbf{L} - \mathbf{l}$, and $\mathbf{l}'' \equiv \mathbf{L}_1 - \mathbf{l}$. Here, $\bar{T} = T$, $\bar{E} = E$ and $\bar{B} = E$, and the geometric functions, h , come from the response function definitions in table 1. In eqs. (6.4) and (6.5) we follow the approximate non-perturbative result of ref. [39] by replacing unlensed CMB spectra with the corresponding lensed gradient spectra, $C^{XY} \rightarrow \tilde{C}^{X\nabla Y}$. We obtain the template rotation bispectra by contraction of the lensing fields with eq. (4.4), giving

$$\begin{aligned} b_{(-L)\mathbf{L}_1\mathbf{L}_2}^{\hat{\omega}^{\text{tem}}\phi\phi} &= \frac{4}{L_1^2 L_2^2} b_{(-L)\mathbf{L}_1\mathbf{L}_2}^{\hat{\omega}^{\text{tem}}\kappa\kappa} \\ &\equiv \frac{4F_L^{-1}}{L_1^2 L_2^2} b_{(-L)\mathbf{L}_1\mathbf{L}_2}^{\omega^{ij}} (\mathbf{C}_{\text{LSS}}^{-1})_{L_1}^{ip} (\mathbf{C}_{\text{LSS}}^{-1})_{L_2}^{jq} C_{L_1}^{p\kappa} C_{L_2}^{q\kappa}. \end{aligned} \quad (6.6)$$

With a non-zero fiducial $b^{\omega^{ij}}$ in the weights, the template-convergence bispectra are non-zero even if the data rotation is zero and all the fields are Gaussian.

Figure 9 shows the total bias from all the $N^{(2)}$ terms for different rotation templates at SO goal sensitivity. Almost all of the analytic predictions for the bias are smaller than the

⁶We postpone discussion till section 6.2 of additional contributions that arise when one or both of \hat{a} , \hat{b} are also derived from a quadratic estimator.

⁷Following the literature convention to label lensing bias terms by their order in $C^{\phi\phi}$ about the Gaussian result, which is equivalent to labelling by $\mathcal{O}(\Psi^2)$.




Figure 9. Analytic predictions for the $N_{A1}^{(2)} + N_{C1}^{(2)}$ bias (dashed lines) on the rotation-template cross-spectrum $C_L^{\hat{\omega}^{\text{tem}}}$ for different templates at SO goal sensitivity, assuming Gaussian fields. The lensing observables, $\hat{\omega}$ and $\hat{\kappa}$, are reconstructed with the GMV estimator. The curves can be compared to the fiducial post-Born rotation power spectrum in solid black, and the null dotted line. Contributions from the contractions entering the $\hat{\kappa}$ quadratic estimator are neglected here, apart from the linear κ field itself.

expected fiducial rotation signal at all multipoles. The smallest biases across most scales are for the optimal template, $\hat{\omega}^{\{\kappa,g,l\}}$, and the convergence-free template, $\hat{\omega}^{\{g,l\}}$. Here we have not included the bias from $\hat{\omega}^{\{\kappa,l\}}$, as the result is around an order of magnitude larger than the other configurations. As shown in section 5, $\hat{\omega}^{\{\kappa,l\}}$ is the noisiest rotation template, and hence the least interesting to consider.

6.2 Additional biases with $\hat{\kappa}$

The $N^{(2)}$ bias terms eq. (6.4) and eq. (6.5), represent the leading biases only under the assumption that \hat{X} and \hat{a} are independent. This is most untrue for $\hat{a} = \hat{\kappa}$, which is reconstructed from pairs of lensed CMB maps, $\hat{\kappa} \propto \tilde{X}\tilde{Y}$. Therefore, the CMB fields used to reconstruct $\hat{\omega}$ will form contractions with the fields used to reconstruct $\hat{\kappa}$, and induce extra biases if $\hat{\kappa}$ is included as a density tracer. There are several configurations in which these additional biases are generated below $N^{(2)}$ level.

Consider the different tracer combinations that go into the construction of $\hat{\omega}^{\text{tem}}$. If $\hat{\kappa}$ is used as one of the input tracers, there will be a term in eq. (4.4) such that $\hat{a}^i \hat{a}^j = \hat{\kappa} \hat{\kappa}$. For this term, the cross-spectrum is related to a 6-point function of CMB maps $\langle \hat{\omega} \hat{\omega}^{\text{tem}} \rangle \propto \langle \tilde{X}\tilde{Y}\tilde{X}\tilde{Y}\tilde{X}\tilde{Y} \rangle$. This has a fully disconnected contribution at $N^{(0)}$ level, leaving 8 terms after Wick expansion, the combination of which produce the total $\hat{\kappa}\hat{\kappa}$ leading order bias

$$N_{\hat{\kappa}\hat{\kappa}}^{(0)}(L) = 2L^2 A_L^\Omega \int \frac{d^2 l d^2 \mathbf{L}_1}{(2\pi)^4} \beta(\mathbf{L}, \mathbf{L}_1) g_{ij}^\Omega(\mathbf{l}, \mathbf{L}) g_{pq}^\phi(\mathbf{l}, \mathbf{L}_1) g_{rs}^\phi(\mathbf{l}', \mathbf{L}_2) (\tilde{C}_{\text{CMB}})_l^{ip} (\tilde{C}_{\text{CMB}})_{l'}^{jr} (\tilde{C}_{\text{CMB}})_{l''}^{qs}, \quad (6.7)$$

XY	$f_{XY}^\phi(\mathbf{l}, \mathbf{L})$
TT	$\mathbf{l} \cdot \mathbf{L} \tilde{C}_l^{T\nabla T} + \mathbf{l}' \cdot \mathbf{L} \tilde{C}_{l'}^{T\nabla T}$
EE	$(\mathbf{l} \cdot \mathbf{L} \tilde{C}_l^{E\nabla E} + \mathbf{l}' \cdot \mathbf{L} \tilde{C}_{l'}^{E\nabla E}) h_E(\mathbf{l}, \mathbf{l}')$
EB	$(\mathbf{l} \cdot \mathbf{L} \tilde{C}_l^{E\nabla E} + \mathbf{l}' \cdot \mathbf{L} \tilde{C}_{l'}^{B\nabla B}) h_B(\mathbf{l}, \mathbf{l}')$
TE	$\mathbf{l} \cdot \mathbf{L} \tilde{C}_l^{T\nabla E} h_E(\mathbf{l}, \mathbf{l}') + \mathbf{l}' \cdot \mathbf{L} \tilde{C}_{l'}^{T\nabla E}$
TB	$\mathbf{l} \cdot \mathbf{L} \tilde{C}_l^{T\nabla E} h_B(\mathbf{l}, \mathbf{l}')$

Table 5. The lensing potential response functions for individual pairs of CMB maps X and Y required for lensing reconstruction. The geometric factors are defined as $h_E(\mathbf{l}_1, \mathbf{l}_2) = \cos(2(\theta_{l_1} - \theta_{l_2}))$, and $h_B(\mathbf{l}_1, \mathbf{l}_2) = \sin(2(\theta_{l_1} - \theta_{l_2}))$, and $\mathbf{l}' = \mathbf{L} - \mathbf{l}$. The non-perturbative response functions given here (slightly generalizing previous results [23, 25]) use the lensed gradient spectra defined in ref. [26]. The BB response function is excluded along with the curl-like terms, \tilde{C}^{XP_\perp} , as they are subdominant relative to the other terms [27].

where the LSS mode coupling function for $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$ is defined by

$$\beta(\mathbf{L}, \mathbf{L}_1) \equiv \frac{L_1^2 L_2^2}{4} A_{L_1}^\phi A_{L_2}^\phi F_L^{-1} b_{(-\mathbf{L})L_1 L_2}^{\omega ij} (\mathbf{C}_{\text{LSS}}^{-1})_{L_1}^{i\kappa} (\mathbf{C}_{\text{LSS}}^{-1})_{L_2}^{j\kappa}. \quad (6.8)$$

Now consider all the different tracer pairs in which $\hat{\kappa}$ accounts for only one leg of the template estimator, $\hat{a}^i \hat{a}^j \in \{\hat{\kappa}g, \hat{\kappa}I, g\hat{\kappa}, I\hat{\kappa}\}$. For each of these cases, the cross-spectrum is now a CMB-tracer 5-point function $\langle \hat{\omega} \hat{\omega}^{\text{tem}} \rangle \propto \langle \tilde{X} \tilde{Y} \tilde{X} \tilde{Y} \hat{a}^j \rangle$. The leading bias in this case is at $N^{(1)}$ level, and is best captured non-perturbatively using definition of the lensing potential response functions (equivalent to eq. (3.3) for the Ω responses; explicit f^ϕ functions given in table 5). There are then 4 distinct terms that contribute to the one- $\hat{\kappa}$ bias at leading order. Using the symmetry property of the GMV weights, $g_{ij}(\mathbf{l}, \mathbf{L}) = g_{ji}(\mathbf{L} - \mathbf{l}, \mathbf{L})$, this bias reduces to

$$N_{\hat{\kappa}}^{(1)}(L) = 2L^2 A_L^\Omega \int \frac{d^2 \mathbf{l} d^2 \mathbf{L}_1}{(2\pi)^4} \alpha(\mathbf{L}, \mathbf{L}_1) g_{ij}^\Omega(\mathbf{l}, \mathbf{L}) g_{pq}^\phi(\mathbf{l}, \mathbf{L}_1) (\tilde{\mathbf{C}}_{\text{CMB}})_{l'}^{ip} f_{jq}^\phi(\mathbf{l}', \mathbf{L}_2), \quad (6.9)$$

complete with definition of the corresponding LSS function

$$\alpha(\mathbf{L}, \mathbf{L}_1) \equiv \frac{L_1^2}{L_2^2} A_{L_1}^\phi F_L^{-1} b_{(-\mathbf{L})L_1 L_2}^{\omega ij} (\mathbf{C}_{\text{LSS}}^{-1})_{L_1}^{i\kappa} (\mathbf{C}_{\text{LSS}}^{-1})_{L_2}^{jq} C_{L_2}^{\kappa q}. \quad (6.10)$$

Figure 10 shows the biases on the rotation cross-spectrum with the optimal template $\hat{\omega}^{\{\kappa, g, I\}}$ at SO goal sensitivity. The $N_{\hat{\kappa}}^{(1)}$ term dominates and is of comparable magnitude to the signal over the multipoles of interest, $500 \lesssim L \lesssim 2000$. The observed CMB power spectra $\tilde{\mathbf{C}}_{\text{CMB}}$ in eq. (6.9) determine the size of this bias. The bias can therefore be reduced by using CMB measurements with no noise correlations, for example if $\hat{\omega}$ and $\hat{\kappa}$ are reconstructed with different sets of CMB maps so that there are no noise correlations between maps used for $\hat{\omega}$ and maps used for $\hat{\kappa}$. The dash-dot lines in figure 10 show the size and shape of $N_{\hat{\kappa}}^{(1)}$ is significantly reduced if the maps have no noise; the remaining $N_{\hat{\kappa}}^{(1)}$ is then comparable in magnitude to $N^{(2)}$. Similarly, $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$ moves closer to the null line on all scales for this case.




Figure 10. Analytic predictions for all biases considered (dashed lines) on the rotation-template cross-spectrum $C_L^{\hat{\omega}^{\{\kappa, g, l\}}}$ at SO goal sensitivity, assuming Gaussian tracer and CMB fields. The lensing observables, $\hat{\omega}$ and $\hat{\kappa}$, are reconstructed with the GMV estimator. The curves can be compared to the fiducial post-Born rotation power spectrum in solid black, and the null dotted line. Contributions from contractions involving the $\hat{\kappa}$ quadratic estimator are shown by the green and orange lines for $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$ and $N_{\hat{\kappa}}^{(1)}$ respectively. If $\hat{\omega}$ and $\hat{\kappa}$ are measured from independent CMB maps then $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$ and $N_{\hat{\kappa}}^{(1)}$ lose the contribution from CMB map level noise and reduce to the corresponding dash-dotted lines.

There are several possible ways to mitigate the contribution of these additional $\hat{\kappa}$ bias signals. One could use $\hat{\kappa}$ and $\hat{\omega}$ reconstructed from maps split up over different time periods [8, 43] to get rid of the noise contributions and reduce the biases back to dash-dot lines in figure 10 (in the approximation that any splits do not affect map level noise). Equivalently, they could be reconstructed from different surveys or detectors. Alternatively, $\hat{\kappa}$ could be reconstructed from T only, while $\hat{\omega}$ is reconstructed with a polarization-only estimator, then only the \tilde{C}_l^{TE} configuration in $(\tilde{C}_{\text{CMB}})_l^{XY}$ survives in eq. (6.10), potentially resulting in another reduction to the bias magnitude. To avoid the $N^{(0)}$ or $N^{(1)}$ entirely, one could simply construct a convergence-free $\hat{\omega}^{\text{tem}}$, for example by only using $\hat{a} \in \{\hat{g}, \hat{I}\}$. We have already shown that such a template is detectable with high significance in the near future, and its bias is small compared to the fiducial rotation signal (figure 9).

We have not included biases for lensing reconstruction with iterative estimators in our investigation, as the analytical form of iterative bias terms requires careful consideration. Additionally, the biases considered here are for full GMV reconstruction only, and do not directly apply to the CMB-S4 deep experimental setup which uses polarization-only reconstruction. Hence, analysis of biases appropriate to CMB-S4 deep is left to future work. Possible bias contributions from correlations between CMB foreground residuals and the LSS tracers have also been neglected. Finally, it is worth re-emphasising that only Gaussian biases are considered here, the investigation of contamination from higher order statistics is left to future work.

6.3 Testing using simulations

We use Gaussian Monte Carlo simulations to test the analytically-derived bias terms from the previous two sections. In this section, all lensing reconstruction is achieved via the full GMV estimator for SO goal and CMB-S4 like experiments, and only the optimal rotation template $\hat{\omega}^{\{\kappa,g,I\}}$ is considered.

For no fiducial rotation in the simulations, the cross spectrum is purely bias:

$$\langle C^{\hat{\omega}\hat{\omega}^{\text{tem}}} \rangle_{\omega=0} = \text{bias} = N^{(0)} + N^{(1)} + N^{(2)} + \mathcal{O}(\Psi^6). \quad (6.11)$$

We therefore reconstruct $\hat{\omega}$ and $\hat{\kappa}$ from flat-sky lensed CMB simulations in which $\omega = 0$ in the maps. The code LENSIT [29] performs the lensing of the CMB maps and the reconstruction. We use the same experimental configurations as the Fisher forecasts, and the maps are again cut at $30 \leq L \leq 3000$ for T , and $30 \leq L \leq 5000$ for E and B to maintain consistency.

In preparation for the template estimator, a κ realization is used to generate Gaussian maps of \hat{g} and \hat{I} via the Cholesky decomposition of their covariances, defined by eq. (2.8), eq. (4.9), and eq. (4.10). This provides the necessary cross-correlations between tracers, and dependence between $\hat{\omega}$ and the tracer maps. The optimal rotation template map, $\hat{\omega}^{\{\kappa,g,I\}}$, is then constructed using the fast configuration-space method of eq. (4.14), with the generated tracer maps plus the lensing reconstructed $\hat{\kappa}$ as inputs.

All maps are generated to have sky coverage of $\sim 41,000 \text{ deg}^2$ corresponding to an $f_{\text{sky}} = 1$. There are $(4096)^2$ pixels in each sky patch, so there are modes resolved up to $L \sim 5100$. We apply cuts on the tracer maps at $30 \leq L \leq 3000$ (with the exception of the CIB which has the same stringent cuts of $110 \leq L \leq 2000$) and all modes within this range are fully resolved. We generate 100 sets of simulations.

We first consider the case of independent CMB noise realizations in the $\hat{\omega}$ and $\hat{\kappa}$ reconstructions, so that no contributions to $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$ or $N_{\hat{\kappa}}^{(1)}$ come from map-level noise ($N_{A1}^{(2)}$ and $N_{C1}^{(2)}$ remain unaffected by the map level noise). Therefore, we do not add noise to the CMB or LSS maps as it would only increase the scatter of the results.⁸ However, noise power is included in the filters used for the GMV reconstruction, i.e. in the inverse covariance weights of eq. (3.2), and the inverse filtered LSS maps in eq. (4.4). The total biases present in these noiseless simulations are represented by the points in right panel of figure 11. The bottom panel shows the residuals between simulation and the analytic prediction, normalized by the fiducial post-Born rotation spectrum. For SO at goal sensitivity there is good qualitative agreement that the shape and size of the total bias is almost entirely described by the analytic predictions. This contrasts with CMB-S4; while there is agreement with the overall shape of the bias, there is an offset where the analytic prediction overestimates the size of the bias.

We also test the total bias with map level noise contributions included. In this case, noise is added to the CMB maps used in the reconstruction of $\hat{\omega}$ and $\hat{\kappa}$. Here we also add noise to the tracer maps. The inclusion of tracer noise does not generate new biases; we have consistently assumed noise is independent between LSS fields throughout. However, it does increase the error bars on the simulated bias signal. This enables comparison between the scatter on the simulated bias with our forecast $1\text{-}\sigma$ error on the rotation cross-spectrum in figure 8 to check whether the variance on the bias affects the Fisher forecasts.

The left panel of figure 11 shows the total simulated bias on the rotation cross-spectrum. Again, the shape and size of the total bias in the SO goal setup is well described by the

⁸Note the $\hat{\omega}$ and $\hat{\kappa}$ maps still have reconstruction noise from CMB fluctuations.

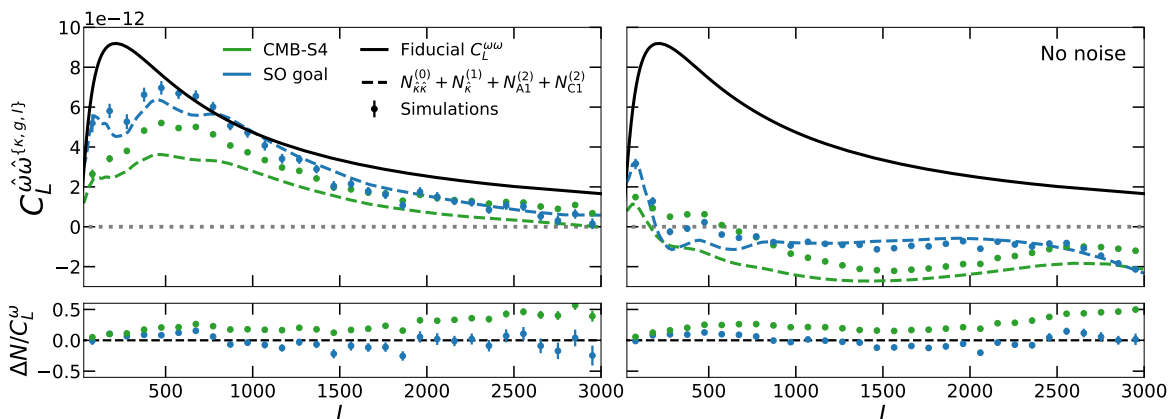


Figure 11. The total Gaussian bias on the optimal rotation-template cross-spectrum $C_L^{\hat{\omega}\hat{\omega}^{\{\kappa,g,l\}}}$ for SO at goal sensitivity (blue) and CMB-S4 (green), assuming $\hat{\omega}$ and $\hat{\kappa}$ are reconstructed with the GMV estimator. Dashed lines show the combined analytic predictions for $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$, $N_{\hat{\kappa}}^{(1)}$, $N_{A1}^{(2)}$ and $N_{C1}^{(2)}$. Points show the mean bias from 100 Monte Carlo simulations for 30 linearly-spaced bins in L , with error bars representing standard error in the mean. The results can be compared to the fiducial post-Born rotation power spectrum in solid black, and the null dotted line. *Left panel:* contributions to the bias from CMB map level noise are included in the theoretical curves and simulated results. This represents the case in which $\hat{\omega}$ and $\hat{\kappa}$ are reconstructed from the same CMB maps. *Right panel:* same as left panel with exception that $\hat{\omega}$ and $\hat{\kappa}$ are now measured using CMB maps with independent noise. Therefore noise is not included in the simulations at map level, and similarly noise contributions are not included in the theory lines. *Bottom panels:* residuals of the total simulated Gaussian bias normalized by fiducial post-Born $C_L^{\omega\omega}$ show qualitative agreement between theory and simulations for SO goal. Unaccounted contributions from $N_{\hat{\kappa}}^{(2)}$ explain the CMB-S4 discrepancy.

analytic predictions. However, the CMB-S4 analytic biases under-predict the total bias from the simulations, though the shapes are qualitatively consistent. The mischaracterization of CMB-S4 biases in both cases considered would constitute a $\sim 20\text{--}50\%$ residual offset on measurement of the fiducial spectrum if reliant purely on the analytic predictions of $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$, $N_{\hat{\kappa}}^{(1)}$, $N_{A1}^{(2)}$ and $N_{C1}^{(2)}$. We find the discrepancy to be consistent with a missing theoretical bias contribution from $N_{\hat{\kappa}}^{(2)}$ like terms which are perhaps more prevalent for CMB-S4 due to lower noise on the polarization. This is explored further in appendix C along with further validation tests for individual bias terms. We leave explicit derivation of $N_{\hat{\kappa}}^{(2)}$ to future work.

Finally, we find the simulations to be in excellent agreement with the size of the forecast error-bars on the rotation cross-spectrum in figure 8 despite inclusion of the biases shown in figure 11. Hence, for near-future CMB experiments, provided any bias can be reasonably estimated from simulations, subtraction of these Gaussian biases from the measured cross-spectrum is sufficient to probe post-Born rotation power without compromising S/N .

7 Conclusions

We have shown lensing rotation to be detectable at high significance by combining a lensing reconstruction measurement with a template constructed from the lensing convergence and/or large-scale structure tracers. Detection of this rotation signal will give a powerful internal consistency check of the cosmological model, data modelling and systematics. Since the signal

is detectable at high significance, the rotation power spectrum shape can also be recovered at some level.

The quadratic template for lensing rotation that we have constructed could also be used to ‘rotation-clean’ the lensing reconstruction signal, allowing use of the cleaned signal as a more stringent null test or probe of non-standard cosmological signals. E-mode polarization observations combined with the rotation template could also be used to remove the rotation-induced B-mode polarization, in a similar way that E modes combined with large-scale structure tracers and lensing convergence can be used to delens the observed B-mode polarization produced by standard gradient lensing. This can potentially improve the sensitivity to B modes from primordial gravitational waves, though the rotation signal only becomes marginally significant for tensor-scalar ratios $r < 10^{-5}$ [20].

We have made a number of simplifying assumptions in this paper, that would have to be relaxed for application to more-realistic data. Although we have tested consistency with simple null-hypothesis simulations to check the size of the leading biases, future work should use more realistic non-Gaussian post-Born lensing simulations such as those in ref. [44]. This could test for correct recovery of a non-zero rotation signal, and quantify the importance of any additional biases. Future work could also quantify the biases that arise when using iterative lensing reconstruction. If necessary, the biases can be reduced by using semi-independent subsets of the data, and by using different subsets it should also be possible to have multiple consistency checks between results from different rotation estimators.

The parity-odd nature of the rotation field makes it relatively insensitive to many systematics that plague other large-scale structure observations, however any systematics that can produce rotation-like effects would have to be assessed carefully, and additional real-life uncertainties could lead to larger variance. Handling uncertainties in the modelling of bias, and other effects relating observations to the underlying matter field, should also be investigated more thoroughly.

Acknowledgments

We thank Julien Carron for help with his PLANCKLENS and LENSIT codes, and further thank Julien and Giulio Fabbian for insightful discussions. We also thank Toshiya Namikawa and Byeonghee Yu for providing details of their CIB modelling. MR is supported by a U.K. Science and Technology Facilities Council (STFC) studentship. AL was supported by the U.K. STFC grants ST/T000473/1 and ST/X001040/1. Parts of this research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory.

A Iterative estimators

The lensing reconstruction noise from iterative lensing estimation can be estimated quite easily from a modification of the standard analytic result for N_0 [29, 45, 46]. At each step of the iteration, the reconstruction is effectively performed on a sky that has been delensed using the current best estimator for the lensing field (which is always dominated by convergence, since the rotation is very small). Instead of using the noise and lensed CMB power spectra to calculate the disconnected Gaussian N_0 variance, the iterative reconstruction noise can be evaluated by using partially delensed CMB power spectra instead. This reduces the reconstruction noise, primarily due to the reduced variance from the lensing-dominated B modes.

The amount of delensing that is possible at each L depends on the cross-correlation of the current lensing estimate $\hat{\phi}_i$ with the truth:

$$\rho_i^2 \equiv \frac{(C^{\phi\hat{\phi}_i})^2}{C^{\hat{\phi}_i\hat{\phi}_i}C^{\phi\phi}} = \frac{C^{\phi\phi}}{C^{\hat{\phi}_i\hat{\phi}_i}} = \frac{C^{\phi\phi}}{C^{\phi\phi} + N_0^{(i)}}, \quad (\text{A.1})$$

where we have assumed the estimator is unbiased in the sense that $C^{\hat{\phi}_i\phi} = C^{\phi\phi}$ and neglected $N^{(1)}$. The lensing remaining after partial delensing has power spectrum $(1 - \rho_i^2)C^{\phi\phi}$, which is what is used to calculate the partially delensed spectra. As iterations progress, the delensed power, and hence $N_0^{(i)}$, converge. This noise estimation procedure differs slightly from the method usually used to estimate delensed B-mode noise using Wiener-filtered E modes [45] (also different from ref. [47]), but in practice the forecast results are very similar.

In addition to internal delensing at each step, it is also possible to use one or more external tracers [45, 48]. The optimal combination of the internal lensing estimator at each step, and an external estimate with cross-correlation, ρ_{ext} , to the true lensing field, has a total cross-correlation given by [48]

$$\rho_{\text{tot}}^2 = \frac{(1 - \rho_{\text{ext}}^2)C^{\phi\phi} + \rho_{\text{ext}}^2 N_0}{(1 - \rho_{\text{ext}}^2)C^{\phi\phi} + N_0}. \quad (\text{A.2})$$

The joint iterative $N_0^{(i)}$ is then estimated simply by partially delensing using this combined correlation function.

B Density tracers

B.1 CMB lensing convergence

The lensing convergence was introduced in section 2.1, and described generally in the wider context of weak lensing. Here κ refers to the CMB lensing convergence only.

The only noise term considered is the lowest order reconstruction bias $N_L^\kappa = N_0^\kappa(L)$ for either the GMV estimator (similar to the terms described in section 3.1 in the context of curl reconstruction), or iterative reconstruction (see appendix A). These noise levels for $C_L^{\kappa\kappa}$ are illustrated in figure 12 for the different experimental setups considered in this paper.

The window function for the convergence is written explicitly in eq. (2.10).

B.2 Galaxy clustering

Galaxies are a natural tracer of the matter distribution. They form within over-densities and are therefore a biased tracer of the LSS. Their window function is

$$W_g(\chi) = b(z)n(z)H(z), \quad (\text{B.1})$$

for bias, b , and Hubble parameter, H , evaluated at redshift $z(\chi)$. The Hubble parameter is required to account for the change in integration variable from z to χ in eq. (4.8). The redshift distribution is normalized as follows

$$n(z) = \frac{dN/dz}{\int dz'(dN/dz')}. \quad (\text{B.2})$$

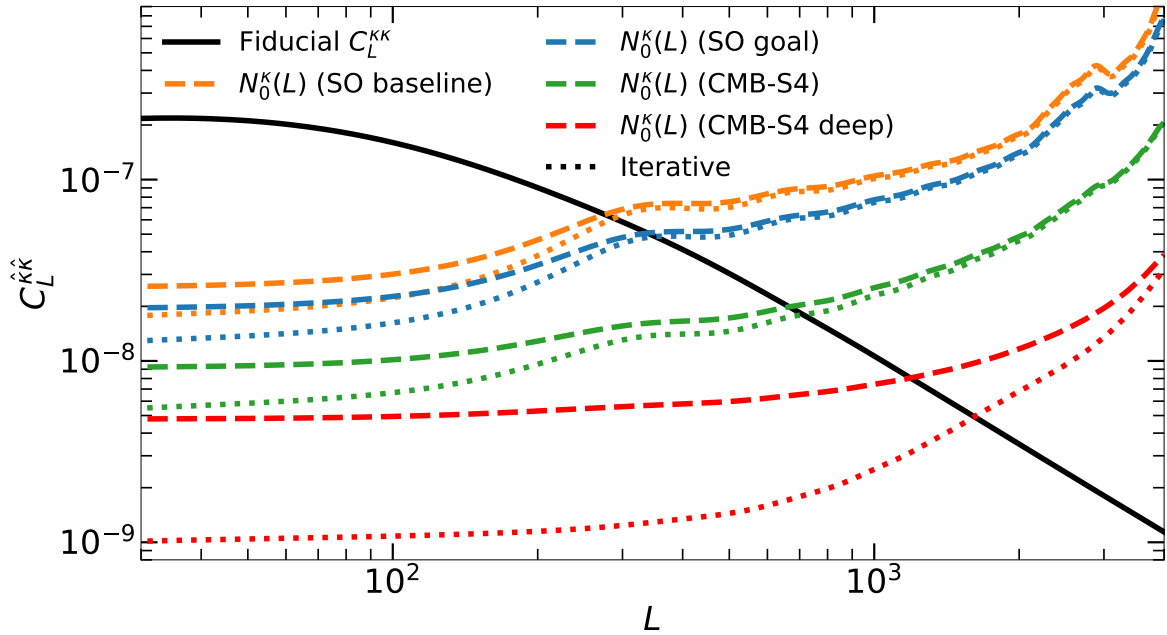


Figure 12. The lowest-order CMB lensing reconstruction noise, N_0^{κ} , for GMV estimators are shown by the dashed lines for SO and CMB-S4 experimental configurations. Iterative reconstruction is shown in the dotted lines for the same experiments. CMB-S4 deep patch uses polarization-only reconstruction, and covers just 5% of the sky and assumes no foregrounds. The fiducial lensing convergence power spectrum is shown in solid black.

We do not split the distribution into redshift bins, instead g is considered for one effective redshift bin described by the LSST gold model [36], which for a survey of limiting magnitude at 25.3 corresponding to $z_0 = 0.311$, is described by

$$\frac{dN}{dz} = \frac{1}{2z_0} \left(\frac{z}{z_0}\right)^2 \exp\left(\frac{-z}{z_0}\right). \quad (\text{B.3})$$

The bias is modelled to linear order, $b(z) = 1 + 0.84z$. We take a conservative forecast for the LSST gold sample as having a mean galaxy count of $\bar{n} = 40$ galaxies/arcmin², and correspondingly the flat shot noise for galaxy counts is

$$N_L^g = N_{\text{shot}}^g = 1/\bar{n}, \quad (\text{B.4})$$

where \bar{n} is converted to units of galaxies per steradian. The total galaxy power spectrum for the LSST gold sample is shown in figure 13.

For simplicity, we do not consider redshift space distortions, magnification bias, or any other subdominant contributions to the galaxy clustering model.

B.3 Cosmic Infrared Background

The CIB radiation is primarily emission from dust in galaxies. Thus, similar to the galaxy density observable, it is also a biased tracer of the LSS. Using the single spectral energy distribution (SED) model [49], the window function at frequency, ν , is defined as

$$W_I^\nu(\chi) = A_\nu \frac{\chi^2}{(1+z(\chi))^2} \exp\left(-\frac{(z-z_c)^2}{2\sigma_z^2}\right) f_{\nu(1+z)}, \quad (\text{B.5})$$




Figure 13. Fiducial galaxy density power spectrum (solid line) and the corresponding shot noise for an LSST-like survey (dashed line).

for

$$f_\nu = \begin{cases} \left[\exp\left(\frac{h\nu}{k_B T}\right) - 1 \right]^{-1} \nu^{\beta+3} & (\nu \leq \nu') \\ \left[\exp\left(\frac{h\nu}{k_B T}\right) - 1 \right]^{-1} \nu^{\beta+3} \left(\frac{\nu}{\nu'}\right)^{-\alpha} & (\nu > \nu'), \end{cases} \quad (\text{B.6})$$

with $z_c = 2$, $\sigma_z = 2$, $\beta = 2$, and $\alpha = 2$.

There are two main components to the noise for the CIB

$$N_L^{I,\nu} = N_{\text{shot}}^{I,\nu} + N_{\text{dust}}^{I,\nu}(L), \quad (\text{B.7})$$

from Poisson noise, and contamination from galactic dust emission. This assumes the CMB power spectra has been accurately removed, a safe assumption given the precision of recent CMB experiments [50]. We do not account for instrumental noise which is expected to dominate on small scales. Instead, we set $L_{\text{max}} = 2000$ for C_L^{II} . Similarly, due to the difficulty in optimally removing galactic dust contamination from the CIB maps, noise from dust dominates on large scales. We set $L_{\text{min}} = 110$ to account for this. Above L_{min} , the dust is modelled by a simple two parameter power law

$$N_{\text{dust}}^{I,\nu}(L) = \beta_\nu L^{-\alpha_\nu}. \quad (\text{B.8})$$

Following [40, 41], the Generalized Needlet Internal Linear Combination (GNILC) algorithm [51] is applied to Planck data [52], and $C_L^{II,\nu} + N_L^{I,\nu}$ is fit to the filtered map to obtain values for A_ν , β_ν , and α_ν .

The total CIB power spectrum at $\nu = 353$ GHz is shown in figure 14 along with the noise contributions from sampling and dust for the multipole range $L_{\text{min}} < L < L_{\text{max}}$.




Figure 14. The main noise contributions to Planck-like Cosmic Infrared Background (CIB) measurements are shown compared to the best-fit spectrum at 353 GHz (solid line). Shot noise (dashed line) and dust noise (dotted line) are individually plotted alongside the total noise (dot-dashed line). All curves here have units of $\text{MJy}^2 \text{sr}^{-1}$.

C Isolating cross-spectrum biases

Here, Gaussian Monte Carlo simulations are used to test the analytically derived results for each source of bias considered in this paper.

C.1 $N_{\text{A1}}^{(2)} + N_{\text{C1}}^{(2)}$

The $N^{(2)}$ biases are tested with Monte Carlo simulations using slight modifications to the methodology presented in section 6.3. The first modification is not adding noise to any of the CMB or LSS maps. Noise does not contribute to either $N_{\text{A1}}^{(2)}$ or $N_{\text{C1}}^{(2)}$; including noise would only increase the scatter of the results. However, noise power is included in the filters used for the GMV reconstruction, i.e. in the inverse covariance weights of eq. (3.2), and the inverse filtered LSS maps in eq. (4.4).

The other modification is not using a lensing reconstructed convergence as one of the LSS tracer maps. Instead, we directly take the input κ realization that performs the lensing on the LENSIT CMB maps as one of the inputs to the template QE. This ensures that the additional bias terms described in section 6.2 from contractions within the quadratic estimation of $\hat{\kappa}$ do not exist within the simulations. This methodology does not actually isolate $N_{\text{A1}}^{(2)}$ or $N_{\text{C1}}^{(2)}$, instead it captures the total Gaussian bias in the absence of these additional $\hat{\kappa}$ induced terms. This will include contributions from biases such as the $\mathcal{O}(\Psi^6)$ terms from eq. (6.3).

The top panel of figure 15 shows the resulting total bias from 100 simulations overlaid with the analytical predictions. The residuals normalized with respect to the $1\text{-}\sigma$ error in the bottom panel demonstrate good agreement that the total bias in this setup is well described by $N_{\text{A1}}^{(2)} + N_{\text{C1}}^{(2)}$. There is a slight discrepancy with the CMB-S4 results at $L > 1500$, suggesting

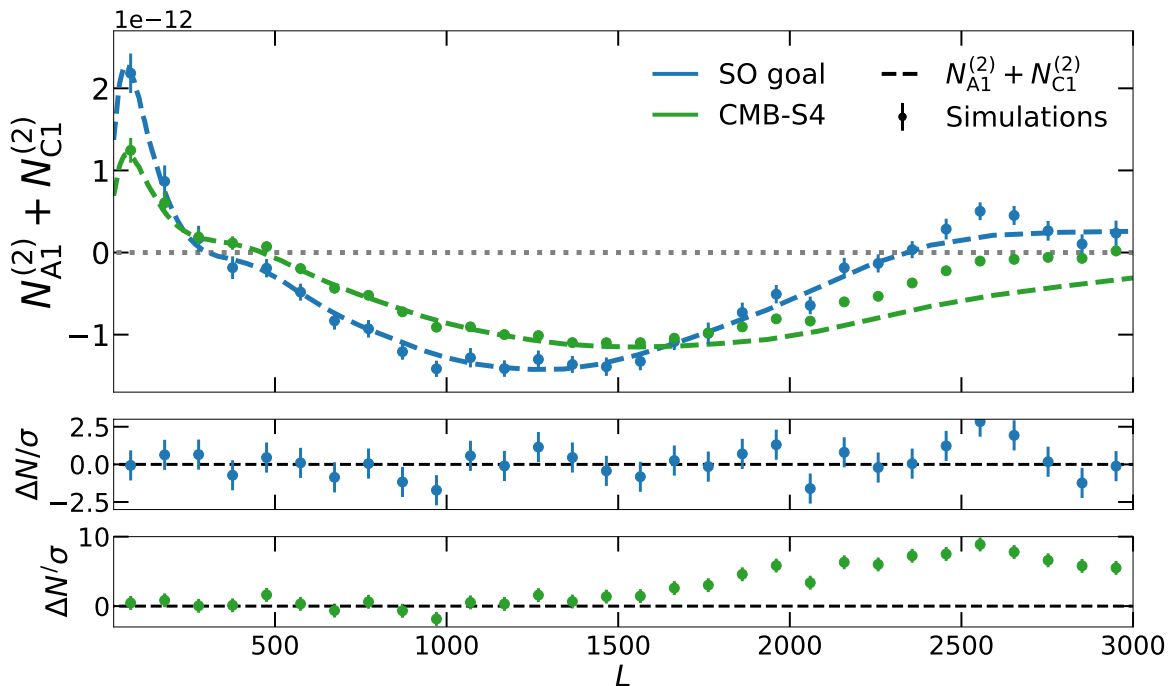


Figure 15. The total Gaussian bias on the optimal rotation-template cross-spectrum $C^{\hat{\omega}\hat{\omega}^{\{\kappa,g,l\}}}$ for SO at goal sensitivity (blue) and CMB-S4 (green), assuming there are no contributions from contractions entering via the $\hat{\kappa}$ quadratic estimator. Therefore, only the combination $N_{A1}^{(2)} + N_{C1}^{(2)}$ (dashed lines) should be present. $\hat{\omega}$ and $\hat{\kappa}$ are reconstructed with the GMV estimator. Points represent the mean bias from 100 Monte Carlo simulations for 30 linearly-spaced bins in L , with error bars representing standard error in the mean. *Bottom panels:* residuals of the total simulated Gaussian bias normalized by their $1\text{-}\sigma$ errors.

possible contributions from biases above $N^{(2)}$ level are getting picked up in the simulations. CMB-S4 would be more susceptible to such terms because of better lensing reconstruction on small scales (due to lower CMB noise, particularly in the polarization maps).

C.2 $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$

To isolate all $N^{(0)}$ bias from simulations we use the same methodology as section 6.3 with one modification. The CMB maps used to “reconstruct” $\hat{\omega}$ and $\hat{\kappa}$ are now Gaussian realizations of lensed CMB power-spectra. As the maps have not been deflected they therefore contain no lensing signal, while retaining the same correlations between maps. The resulting simulated rotation cross-spectrum can only contain bias signals at $N^{(0)}$ level as $\langle\phi a\rangle$ correlations in e.g. eq. (6.6) and eq. (6.10) (or $\langle\phi\phi\rangle$ in higher order terms) vanish.

We carry out the simulations for two cases; the first is with noise included in all maps, and the second is without noise. The results of these two cases are displayed in the left and right panels of figure 16. The total $N^{(0)}$ biases from 100 simulations are overlaid with the analytical predictions for $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$. The residuals normalized with respect to the $1\text{-}\sigma$ error in the bottom panels demonstrate good agreement that the total bias in both setups are well described by $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$.




Figure 16. The total $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$ bias on the optimal rotation-template cross-spectrum $C^{\hat{\omega}\hat{\omega}\{\kappa,g,l\}}$ for SO at goal sensitivity (blue) and CMB-S4 (green), assuming $\hat{\omega}$ and $\hat{\kappa}$ are reconstructed with the GMV estimator. Dashed lines show the analytic predictions for $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$. Points show the mean bias from 100 Monte Carlo simulations for 20 linearly-spaced bins in L , with error bars representing standard error in the mean. *Left panel:* contributions to the bias from CMB map level noise are included in the theoretical curves and simulated results. This represents the case in which $\hat{\omega}$ and $\hat{\kappa}$ are reconstructed from the same CMB maps. *Right panel:* same as left panel with exception that $\hat{\omega}$ and $\hat{\kappa}$ are now measured using CMB maps with independent noise. Therefore noise is not included in the simulations at map level, and similarly noise contributions are not included in the theory lines. *Bottom panels:* residuals of the total simulated Gaussian bias normalized by their $1\text{-}\sigma$ errors show good agreement between theory and simulations.

C.3 $N_{\hat{\kappa}}^{(1)}$

To test $N_{\hat{\kappa}}^{(1)}$ with simulations we follow a similar procedure to refs. [2, 53]. For each set of LSS simulations, we generate 2 independent realizations of unlensed CMB maps, $X^1 \in \{T^1, E^1, B^1\}$ and $X^2 \in \{T^2, E^2, B^2\}$, and use LENSIT to lens each set such that all 6 CMB maps are lensed with the same realization $\phi = \phi_1$. Then the lensing observables are quadratically estimated with the usual GMV estimator (eq. (3.1) for $\hat{\omega}$), but in this case each QE leg contains a map from a different set, e.g. in eq. (3.1) $\tilde{X}^i \in \tilde{X}_{\phi_1}^1$ and $\tilde{X}^j \in \tilde{X}_{\phi_1}^2$. Explicitly, we reconstruct the lensing observables such that

$$\hat{\omega}^{X^1, X^2} \equiv \frac{\hat{\omega} [\tilde{X}_{\phi_1}^1, \tilde{X}_{\phi_1}^2] + \hat{\omega} [\tilde{X}_{\phi_1}^2, \tilde{X}_{\phi_1}^1]}{2}, \quad (\text{C.1})$$

and

$$\hat{\kappa}^{X^1, X^2} \equiv \frac{\hat{\kappa} [\tilde{X}_{\phi_1}^1, \tilde{X}_{\phi_1}^2] + \hat{\kappa} [\tilde{X}_{\phi_1}^2, \tilde{X}_{\phi_1}^1]}{2}. \quad (\text{C.2})$$

The same maps are provided as input to $\hat{\omega}^{X^1, X^2}$ and $\hat{\kappa}^{X^1, X^2}$. This guarantees the independence between pairs of lensed CMB fields within the same QE, hence $N_{A1}^{(2)}$, $N_{C1}^{(2)}$, and $N_{\hat{\kappa}\hat{\kappa}}^{(0)}$ vanish. However, CMB cross-spectrum signals still exist from contractions between the $\hat{\omega}^{X^1, X^2}$ and $\hat{\kappa}^{X^1, X^2}$ estimators, so $N_{\hat{\kappa}}^{(1)}$ like signals still get picked up by the simulation. From

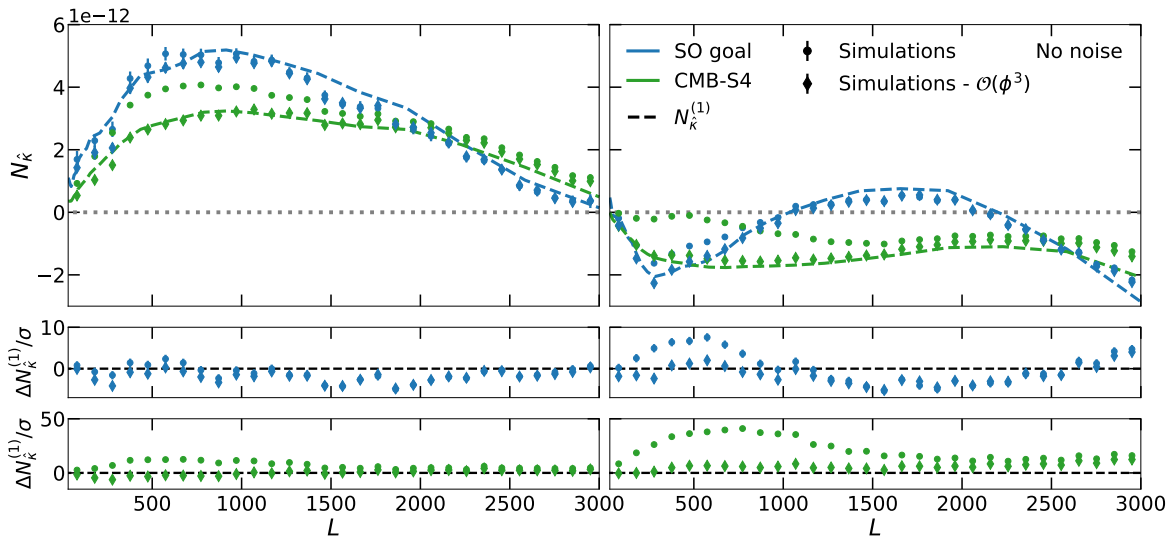


Figure 17. Isolation of the one- κ bias terms, with leading contribution from $N_{\hat{\kappa}}^{(1)}$, on the optimal rotation-template cross-spectrum $C^{\hat{\omega}\hat{\omega}^{\{\kappa,g,I\}}}$ for SO at goal sensitivity (blue) and CMB-S4 (green) assuming $\hat{\omega}$ and $\hat{\kappa}$ are reconstructed with the GMV estimator. Dashed lines show the analytic predictions for $N_{\hat{\kappa}}^{(1)}$. Points show the mean bias from 100 Monte Carlo simulations for 30 linearly-spaced bins in L , with error bars representing standard error in the mean. The diamond points are from the same simulation but with contributions that scale as $N_{\hat{\kappa}}^{(2)}$ subtracted, i.e. subtracting terms that scale as $\mathcal{O}(\phi^3 \times \text{tracer})$ estimated using a scaled simulation. *Left panel:* contributions to the bias from CMB map level noise are included in the theoretical curves and simulated results. This represents the case in which $\hat{\omega}$ and $\hat{\kappa}$ are reconstructed from the same CMB maps. *Right panel:* same as left panel with exception that $\hat{\omega}$ and $\hat{\kappa}$ are now measured using CMB maps with independent noise. Therefore noise is not included in the simulations at map level, and similarly noise contributions are not included in the theory lines. *Bottom panels:* residuals of the total simulated Gaussian bias normalized by their $1\text{-}\sigma$ errors (dots) show qualitative agreement between theory and simulations for SO goal. Residuals of the simulated bias with $N_{\hat{\kappa}}^{(2)}$ removed (diamonds) reduce the CMB-S4 discrepancy, and improve SO goal results at low L .

the definitions eq. (C.1) and eq. (C.2), the simulations under-predict $N_{\hat{\kappa}}^{(1)}$ by a factor of two, this is accounted for by doubling the final simulated output.

Another consequence of this method is the $\hat{a}^i \hat{a}^j = \hat{\kappa} \hat{\kappa}$ term in eq. (4.4) always vanishes, therefore $N_{\hat{\kappa}\hat{\kappa}}^{(1)}$ biases vanish. However, the same is also true for the $N_{\hat{\kappa}\kappa}^{(1)}$ like terms that should be present in eq. (6.9). To account for this, the $N_{\hat{\kappa}}^{(1)}$ predictions in figure 17 do not include the $\hat{\kappa}\hat{\kappa}$ term in the numerical integration of eq. (6.9).

We again perform the simulations for two cases; the first with noise included in all maps, and the second without noise. The results are displayed by the dots in figure 17 for 100 simulations. There is decent agreement between simulation and theory $N_{\hat{\kappa}}^{(1)}$ for SO goal (blue dashed lines) only. The CMB-S4 simulated bias is discrepant compared to the analytic curves (green dashed lines) with and without noise, similar to the total bias results in figure 11.

The disagreement between simulations and $N_{\hat{\kappa}}^{(1)}$ for CMB-S4 hints at missing contributions to the bias. Consider again the arguments for the leading one- κ bias term, we have $\langle \hat{\omega} \hat{\omega}^{\text{tem}} \rangle \propto \langle \tilde{X} \tilde{Y} \tilde{X} \tilde{Y} \hat{a}^j \rangle$. Now expanding out the CMB fields using eq. (6.2) we find next-to-

leading order contractions such as $\langle \delta_\phi X \tilde{Y} \delta_\phi X \delta_\phi Y \hat{a}^j \rangle$, $\langle \delta_\phi^2 X \tilde{Y} \tilde{X} \delta_\phi Y \hat{a}^j \rangle$, $\langle \delta_\phi^3 X \tilde{Y} \tilde{X} \tilde{Y} \hat{a}^j \rangle$, etc. at $N^{(2)}$ level. These terms form part of the $N_{\hat{\kappa}}^{(2)}$ bias, and are not cancelled by the simulated configuration in figure 17. To verify that the discrepancy can be attributed to $N_{\hat{\kappa}}^{(2)}$ we perform a simple scaling test. We re-compute the simulations with $\phi_1 \rightarrow 2\phi_1$, and find that the resulting offset increased by a factor of ~ 8 as expected if $\Delta N \propto N_{\hat{\kappa}}^{(2)}$ (from contractions of three ϕ with one external tracer). The offset is then rescaled and subtracted from the total simulated bias, producing the diamond points in figure 17. This provides better agreement with theory $N_{\hat{\kappa}}^{(1)}$ and also improves SO goal at $L < 1000$. Therefore, consideration of $N_{\hat{\kappa}}^{(2)}$ bias should also resolve the issues in figure 11 for CMB-S4. Derivation and further validation of $N_{\hat{\kappa}}^{(2)}$ is left to future work, but the full total bias can easily be estimated from simulations if the data can be simulated reliably.

References

- [1] A. Lewis and A. Challinor, *Weak gravitational lensing of the CMB*, *Phys. Rept.* **429** (2006) 1 [[astro-ph/0601594](#)] [[INSPIRE](#)].
- [2] PLANCK collaboration, *Planck 2018 results. Part VIII. Gravitational lensing*, *Astron. Astrophys.* **641** (2020) A8 [[arXiv:1807.06210](#)] [[INSPIRE](#)].
- [3] W.L.K. Wu et al., *A Measurement of the Cosmic Microwave Background Lensing Potential and Power Spectrum from 500 deg² of SPTpol Temperature and Polarization Data*, *Astrophys. J.* **884** (2019) 70 [[arXiv:1905.05777](#)] [[INSPIRE](#)].
- [4] B.D. Sherwin et al., *Two-season Atacama Cosmology Telescope polarimeter lensing power spectrum*, *Phys. Rev. D* **95** (2017) 123529 [[arXiv:1611.09753](#)] [[INSPIRE](#)].
- [5] BICEP2 and KECK ARRAY collaborations, *BICEP2/Keck Array. Part VIII. Measurement of gravitational lensing from large-scale B-mode polarization*, *Astrophys. J.* **833** (2016) 228 [[arXiv:1606.01968](#)] [[INSPIRE](#)].
- [6] POLARBEAR collaboration, *Internal delensing of Cosmic Microwave Background polarization B-modes with the POLARBEAR experiment*, *Phys. Rev. Lett.* **124** (2020) 131301 [[arXiv:1909.13832](#)] [[INSPIRE](#)].
- [7] J. Carron, M. Mirmelstein and A. Lewis, *CMB lensing from Planck PR4 maps*, *JCAP* **09** (2022) 039 [[arXiv:2206.07773](#)] [[INSPIRE](#)].
- [8] ACT collaboration, *The Atacama Cosmology Telescope: A Measurement of the DR6 CMB Lensing Power Spectrum and its Implications for Structure Growth*, [arXiv:2304.05202](#) [[INSPIRE](#)].
- [9] B. Jain, U. Seljak and S.D.M. White, *Ray tracing simulations of weak lensing by large scale structure*, *Astrophys. J.* **530** (2000) 547 [[astro-ph/9901191](#)] [[INSPIRE](#)].
- [10] A. Cooray and W. Hu, *Second order corrections to weak lensing by large scale structure*, *Astrophys. J.* **574** (2002) 19 [[astro-ph/0202411](#)] [[INSPIRE](#)].
- [11] E. Krause and C.M. Hirata, *Weak lensing power spectra for precision cosmology: Multiple-deflection, reduced shear and lensing bias corrections*, *Astron. Astrophys.* **523** (2010) A28 [[arXiv:0910.3786](#)] [[INSPIRE](#)].
- [12] G. Pratten and A. Lewis, *Impact of post-Born lensing on the CMB*, *JCAP* **08** (2016) 047 [[arXiv:1605.05662](#)] [[INSPIRE](#)].
- [13] D. Beck, G. Fabbian and J. Errard, *Lensing Reconstruction in Post-Born Cosmic Microwave Background Weak Lensing*, *Phys. Rev. D* **98** (2018) 043512 [[arXiv:1806.01216](#)] [[INSPIRE](#)].

- [14] A. Cooray, M. Kamionkowski and R.R. Caldwell, *Cosmic shear of the microwave background: The Curl diagnostic*, *Phys. Rev. D* **71** (2005) 123527 [[astro-ph/0503002](#)] [[INSPIRE](#)].
- [15] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](#)] [[INSPIRE](#)].
- [16] K. Abazajian et al., *CMB-S4 Science Case, Reference Design, and Project Plan*, [arXiv:1907.04473](#) [[INSPIRE](#)].
- [17] PLANCK collaboration, *Planck 2015 results. Part XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [[arXiv:1502.01589](#)] [[INSPIRE](#)].
- [18] A. Lewis, A. Challinor and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys. J.* **538** (2000) 473 [[astro-ph/9911177](#)] [[INSPIRE](#)].
- [19] A. Mead, J. Peacock, C. Heymans, S. Joudaki and A. Heavens, *An accurate halo model for fitting non-linear cosmological power spectra and baryonic feedback models*, *Mon. Not. Roy. Astron. Soc.* **454** (2015) 1958 [[arXiv:1505.07833](#)] [[INSPIRE](#)].
- [20] A. Lewis, A. Hall and A. Challinor, *Emission-angle and polarization-rotation effects in the lensed CMB*, *JCAP* **08** (2017) 023 [[arXiv:1706.02673](#)] [[INSPIRE](#)].
- [21] N. Kaiser, *Weak gravitational lensing of distant galaxies*, *Astrophys. J.* **388** (1992) 272 [[INSPIRE](#)].
- [22] M. LoVerde and N. Afshordi, *Extended Limber Approximation*, *Phys. Rev. D* **78** (2008) 123506 [[arXiv:0809.5112](#)] [[INSPIRE](#)].
- [23] W. Hu and T. Okamoto, *Mass reconstruction with CMB polarization*, *Astrophys. J.* **574** (2002) 566 [[astro-ph/0111606](#)] [[INSPIRE](#)].
- [24] T. Okamoto and W. Hu, *CMB lensing reconstruction on the full sky*, *Phys. Rev. D* **67** (2003) 083002 [[astro-ph/0301031](#)] [[INSPIRE](#)].
- [25] A.S. Maniyar, Y. Ali-Haïmoud, J. Carron, A. Lewis and M.S. Madhavacheril, *Quadratic estimators for CMB weak lensing*, *Phys. Rev. D* **103** (2021) 083524 [[arXiv:2101.12193](#)] [[INSPIRE](#)].
- [26] A. Lewis, A. Challinor and D. Hanson, *The shape of the CMB lensing bispectrum*, *JCAP* **03** (2011) 018 [[arXiv:1101.2234](#)] [[INSPIRE](#)].
- [27] D. Hanson, A. Challinor, G. Efstathiou and P. Bielewicz, *CMB temperature lensing power reconstruction*, *Phys. Rev. D* **83** (2011) 043005 [[arXiv:1008.4403](#)] [[INSPIRE](#)].
- [28] C.M. Hirata and U. Seljak, *Reconstruction of lensing from the cosmic microwave background polarization*, *Phys. Rev. D* **68** (2003) 083002 [[astro-ph/0306354](#)] [[INSPIRE](#)].
- [29] J. Carron and A. Lewis, *Maximum a posteriori CMB lensing reconstruction*, *Phys. Rev. D* **96** (2017) 063510 [[arXiv:1704.08230](#)] [[INSPIRE](#)].
- [30] L. Knox, *Determination of inflationary observables by cosmic microwave background anisotropy experiments*, *Phys. Rev. D* **52** (1995) 4307 [[astro-ph/9504054](#)] [[INSPIRE](#)].
- [31] R.A. Fisher, *On the Mathematical Foundations of Theoretical Statistics*, *Phil. Trans. Roy. Soc. Lond. A* **222** (1922) 309 [[INSPIRE](#)].
- [32] M. Tegmark, *Measuring cosmological parameters with galaxy surveys*, *Phys. Rev. Lett.* **79** (1997) 3806 [[astro-ph/9706198](#)] [[INSPIRE](#)].
- [33] W. Hu, *Weak lensing of the CMB: A harmonic approach*, *Phys. Rev. D* **62** (2000) 043007 [[astro-ph/0001303](#)] [[INSPIRE](#)].
- [34] W. Hu, *Power spectrum tomography with weak lensing*, *Astrophys. J. Lett.* **522** (1999) L21 [[astro-ph/9904153](#)] [[INSPIRE](#)].

- [35] A. Maraido, A. Hall and A. Taylor, *Testing quadratic maximum likelihood estimators for forthcoming Stage-IV weak lensing surveys*, *Mon. Not. Roy. Astron. Soc.* **520** (2023) 4836 [[arXiv:2207.10412](#)] [[INSPIRE](#)].
- [36] LSST SCIENCE and LSST PROJECT collaborations, *LSST Science Book, Version 2.0*, [arXiv:0912.0201](#) [[INSPIRE](#)].
- [37] D.B. Thomas, L. Whittaker, S. Camera and M.L. Brown, *Estimating the weak-lensing rotation signal in radio cosmic shear surveys*, *Mon. Not. Roy. Astron. Soc.* **470** (2017) 3131 [[arXiv:1612.01533](#)] [[INSPIRE](#)].
- [38] A. Lewis, *The real shape of non-Gaussianities*, *JCAP* **10** (2011) 026 [[arXiv:1107.5431](#)] [[INSPIRE](#)].
- [39] G. Fabbian, A. Lewis and D. Beck, *CMB lensing reconstruction biases in cross-correlation with large-scale structure probes*, *JCAP* **10** (2019) 057 [[arXiv:1906.08760](#)] [[INSPIRE](#)].
- [40] B. Yu, J.C. Hill and B.D. Sherwin, *Multitracer CMB delensing maps from Planck and WISE data*, *Phys. Rev. D* **96** (2017) 123511 [[arXiv:1705.02332](#)] [[INSPIRE](#)].
- [41] T. Namikawa et al., *Simons Observatory: Constraining inflationary gravitational waves with multitracer B-mode delensing*, *Phys. Rev. D* **105** (2022) 023511 [[arXiv:2110.09730](#)] [[INSPIRE](#)].
- [42] V. Böhm, M. Schmittfull and B.D. Sherwin, *Bias to CMB lensing measurements from the bispectrum of large-scale structure*, *Phys. Rev. D* **94** (2016) 043519 [[arXiv:1605.01392](#)] [[INSPIRE](#)].
- [43] M.S. Madhavacheril, K.M. Smith, B.D. Sherwin and S. Naess, *CMB lensing power spectrum estimation without instrument noise bias*, *JCAP* **05** (2021) 028 [[arXiv:2011.02475](#)] [[INSPIRE](#)].
- [44] G. Fabbian, M. Calabrese and C. Carbone, *CMB weak-lensing beyond the Born approximation: a numerical approach*, *JCAP* **02** (2018) 050 [[arXiv:1702.03317](#)] [[INSPIRE](#)].
- [45] K.M. Smith, D. Hanson, M. LoVerde, C.M. Hirata and O. Zahn, *Delensing CMB Polarization with External Datasets*, *JCAP* **06** (2012) 014 [[arXiv:1010.0048](#)] [[INSPIRE](#)].
- [46] L. Legrand and J. Carron, *Lensing power spectrum of the cosmic microwave background with deep polarization experiments*, *Phys. Rev. D* **105** (2022) 123519 [[arXiv:2112.05764](#)] [[INSPIRE](#)].
- [47] S.C. Hotinli, J. Meyers, C. Trendafilova, D. Green and A. van Engelen, *The benefits of CMB delensing*, *JCAP* **04** (2022) 020 [[arXiv:2111.15036](#)] [[INSPIRE](#)].
- [48] B.D. Sherwin and M. Schmittfull, *Delensing the CMB with the Cosmic Infrared Background*, *Phys. Rev. D* **92** (2015) 043005 [[arXiv:1502.05356](#)] [[INSPIRE](#)].
- [49] N.R. Hall et al., *Angular Power Spectra of the Millimeter Wavelength Background Light from Dusty Star-forming Galaxies with the South Pole Telescope*, *Astrophys. J.* **718** (2010) 632 [[arXiv:0912.4315](#)] [[INSPIRE](#)].
- [50] PLANCK collaboration, *Planck 2018 results. Part VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [51] M. Remazeilles, J. Delabrouille and J.-F. Cardoso, *Foreground component separation with generalised ILC*, *Mon. Not. Roy. Astron. Soc.* **418** (2011) 467 [[arXiv:1103.1166](#)] [[INSPIRE](#)].
- [52] PLANCK collaboration, *Planck intermediate results. Part XLVIII. Disentangling Galactic dust emission and cosmic infrared background anisotropies*, *Astron. Astrophys.* **596** (2016) A109 [[arXiv:1605.09387](#)] [[INSPIRE](#)].
- [53] SPT collaboration, *A Measurement of the Cosmic Microwave Background Gravitational Lensing Potential from 100 Square Degrees of SPTpol Data*, *Astrophys. J.* **810** (2015) 50 [[arXiv:1412.4760](#)] [[INSPIRE](#)].